Popularity Signals in Trial-Offer Markets with Social Influence and Position Bias

Felipe Maldonado[*] Pascal Van Hentenryck[†] Gerardo Berbeglia[‡]
Franco Berbeglia[§]

Abstract

This paper considers trial-offer markets where consumer preferences are modeled by a multinomial logit with social influence and position bias. The social signal for a product i is of the form d_i^r, i.e., its cumulative downloads raised to power r. The paper shows that, when $0 < r < 1$ and a static position assignment (e.g., a quality ranking) is used, the market converges to a unique equilibrium where the market shares depend only on product quality, not their initial appeals or the early dynamics. When $r > 1$, the market become unpredictable and goes most likely to a monopoly for some product: Which product becomes a monopoly depends on the initial conditions of the market. These theoretical results are complemented by an agent-based simulation which indicates that convergence is fast when $0 < r < 1$ and that the quality ranking dominates the well-known popularity ranking in terms of market efficiency. These results shed a new light on the role of social influence which is often blamed for unpredictability, inequalities, and inefficiencies in markets. In contrast, this paper shows that, with a proper social signal and position assignment for the products, the market becomes predictable, and inequalities and inefficiencies can be controlled appropriately.

1 Introduction

The impact of social influence and product visibilities on consumer behavior in Trial-Offer (T-O) markets[1] has been explored in a variety of settings (e.g., [20, 24, 27]). Social influence can be dispensed through different types of social signals: A market place may report the number of past purchases of a product, its consumer ratings, and/or its consumer recommendations. Recent studies [10, 27] however came to the conclusion that the popularity signal (i.e., the number of past purchases or the market share) has a much stronger impact on consumer behavior than the average consumer rating signal.[2] These two experimental studies were conducted in very different settings, using the Android application platform in one case and hotel selection in the other. Consumer preferences are also influenced by product visibilities, a phenomenon that has been widely observed in internet advertisement (e.g., [8]), in online stores such as Expedia, Amazon, and iTunes, as well as physical retail stores (see, e.g., [15]).

[*]The Australian National University and DATA61, Canberra, Australia (felipe.maldonado@data61.csiro.au)
[†]Industrial and Operations Engineering, University of Michigan, USA (pvanhent@umich.edu).
[‡]Melbourne Business School, University of Melbourne, Australia (g.berbeglia@mbs.edu)
[§]Tepper Business School, Carnegie Mellon University, USA (fberbegl@andrew.cmu.edu)

[1]A trial-offer market is a setting where consumers can try products before deciding whether to buy them.
[2]The music market iTunes shows the normalized market share of each song of an album.

Despite its widespread use, there is considerable debate in the scientific community about the benefits of social influence. Many researcher have pointed out the potential negative effects of social influence. The seminal work of Salganik et al [20] on the MusicLab experimental market demonstrated that social influence can introduce significant unpredictability, inequality, and inefficiency in T-O markets. These results were reproduced by many researchers (e.g., [18, 21, 22, 25]). More recently, Hu et al [11] studied a newsvendor problem with two substitutable products with the same quality in which consumer preferences are affected by past purchases. The authors showed that the market is unpredictable but can become less so if one of products has an initial advantage. Other researchers have focused on understanding when these undesirable side-effects arise and where they come from. Ceyhan et al [6] studied a market specified by a logit model where a constant J captures the strength of the signal social. They showed that the market behavior (e.g., whether it is predictable) depends on the strength of the social signal. Their results did not consider product visibilities, which is another important aspect of T-O markets. Indeed, various researchers (e.g., [1, 2, 14, 26]) indicated that unpredictability and inefficiencies are often dependent on how products are displayed in the market. In particular, Abeliuk et al [2] show that social influence is always beneficial in expectation when the products are ordered by the *performance ranking* that maximizes the purchases greedily at each step. This result was obtained using the generalized multinomial logit model proposed by Krumme et al [12] to reproduce the MusicLab experiments. Van Hentenryck et al [26] proves a similar result for the *quality ranking* that assigns the highest quality products to the most visible positions. In addition, they show that the market converges to a monopoly for the highest quality product. These results contrast with the MusicLab experiments which relied on the *popularity ranking* that dynamically assigns the most popular products to the most visible positions.

This paper seeks to expand our understanding of social influence in T-O markets and explores the role of the social signal in conjunction with product visibilities. Our starting point is the generalized multinomial logic model of Krumme et al [12], which we extend to vary the strength of the social signal. More precisely, this paper considers a T-O market where the probability of purchasing product i at time t is given by

$$p_i(\phi^t) = \frac{v_{\sigma(i)} q_i (\phi_i^t)^r}{\sum_{j=1}^{n} v_{\sigma(j)} q_j (\phi_j^t)^r} \tag{1}$$

where v_i is the visibility of the product in position i, q_i is the inherent quality of product i, ϕ_i^t is the market share of product i at time t, and $r > 0$ is the strength of the social signal. As should be clear from the discussion above, prior work on T-O markets with poduct visibilities (e.g., [1, 2, 14, 26]) focused on the case of a linear social signal ($r = 1$). The primary objective of this paper is to understand what happens to the T-O market when $r < 1$.

The paper contains both theoretical and simulation results and its contributions can be summarized as follows:

1. When $r < 1$ and a static ranking is used, *the market converges to a unique equilibrium, which we characterize analytically.* In the equilibrium, the market shares depend only on the product qualities q_i and no monopoly occurs. Moreover, a product of higher quality receives a larger market share than a product of lower quality, introducing a notion of fairness in the market and reducing the inequalities introduced by a linear social signal.

2. When $r > 1$ and a static ranking is used, the market converges almost surely to a monopoly for some product: Which product wins the entire market share depends on the initial condition

and the early dynamics. We also show that there exists an equilibrium which is not a monopoly but this equilibrium is not stable.

3. Agent-based simulations show that the market converges quickly towards an equilibrium when using sublinear social signals and the quality ranking. They also show that the quality ranking outperforms the popularity ranking in maximizing the efficiency of the market. The popularity ranking is also shown to have some significant drawbacks in some settings.

These theoretical results indicate that, when the social influence signal is a sublinear function of the market share and a static ranking of the products (e.g., the quality ranking) is used, the market is entirely predictable, depends only on the product quality, and does not lead to a monopoly. This contrasts with the case of $r = 1$ where the market is entirely predictable but goes to a monopoly for the product of highest quality (assuming the quality ranking) [26] and the case of $r > 1$ where the market becomes unpredictable (even with a static ranking). As a result, sublinear social signals provide a way to balance market efficiency and the inequalities introduced by social influence. In particular, with sublinear social signals and a static ranking, markets do not exhibit a Matthew effect where the winner takes all, and remain predictable.

The remaining of this paper is organized as follows. Section 2 describes the related work. Section 3 introduces T-O markets and the generalized multinomial logit model for consumer preferences considered here. Section 4 reviews some necessary mathematical preliminaries, including the fact that T-O markets can be modeled as Robbins-Monro algorithms. Section 5 derives the equilibria for the market as a function of the social signal and also presents the convergence results. Section 6 reports the results from the agent-based simulation. Section 7 discusses some additional results on sublinear signals. Section 8 discusses the results and concludes the paper.

2 Related Work

The research presented in this paper was motivated by the seminal work of Salganik et al [20]. They study an experimental market called the MusicLab, where participants were presented a list of unknown songs from unknown bands, each song being described by its name and band. The participants were partitioned into two groups exposed to two different experimental conditions: the *independent* condition and the *social influence* condition. In the independent group, participants were shown the songs in a random order and they were allowed to listen to a song and then to download it if they so wish. In the second group (social influence condition), participants were shown the songs in popularity order, i.e., by assigning the most popular songs to the most visible positions. Moreover, these participants were also shown a social signal, i.e., the number of times each song was downloaded. In order 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 examplifies a T-O market where each song represents a product, and listening and downloading a song represent trying and purchasing a product respectively. The results in [20] show that the different worlds evolve differently from one another, and significantly so, providing evidence that social influence may introduce unpredictability, inequalities, and inefficiency in the market.

The results in [20] were reproduced by numerous researchers (e.g., [18, 21, 22, 25]) and, in particular, by Krumme et al [12] who model the MusicLab experiment with a generalized multinomial logit where product utilities depend on the song appeal, quality, visibility, and a social influence signal representing past purchases. The T-O market studied in this paper generalizes this model further by exploring various strengths for the social signal as indicated in Equation 1. The case of a linear signal ($r = 1$) has been given significant attention. Abeliuk et al [2] proposed the *performance ranking* which orders the products optimally at each time t given the appeals, qualities, visibilities, and market shares. They show that, when the performance ranking is used, the market always benefits from social influence in expectation. Van Hentenryck et al [26] study the quality ranking which ranks the products by quality: They show that the quality ranking and, more generally, any static ranking, always benefits from social influence in expectation. They also prove that the market converges almost surely to a monopoly for the highest-quality product, indicating that the quality ranking is both optimal and predictable asymptotically. These results extend well-known theorems on Polya urns and their generalizations (e.g., [7, 17, 19]). Abeliuk et al [1] also show that the performance ranking converges to a monopoly for a linear social signal.

Ceyhan et al [6] study a general choice probability $C_i^J(\phi^t)$, where J represents the strength of the social signal, and prove some general convergence results under some assumptions. In particular, they use the ODE method [16] and a stochastic Lyapunov function (e.g., [13]) to prove that the market converges to an equilibrium when the Jacobian of C_i^J is symmetric (which is not the case when product visibilities are present). They also study in detail the case where the market follows a logit model of the form

$$C_i^J(\phi) = \frac{e^{J\phi_i + q_i}}{\sum_j e^{J\phi_j + q_j}}$$

where J is a constant capturing the strength of the social influence signal. They show that there exists a parameter J^* such that the market converges toward a unique equilibrium when $J < J^*$ and to a monopoly when $J \geq J^*$. No analytical characterization of the equilibrium when $J < J^*$ is presented.

It is interesting to contrast these and our results. Observe first that *the proof technique used in [6] relies on the fact that the Jacobian of C_i^J is symmetric, which is not the case for T-O markets with product visibilities.* Our paper studies such T-O markets and show that, when $0 < r < 1$ and a static ranking is used, *the market converges to an inner equilibrium, which we characterize analytically.* When $r = 1$, the T-O market converges to a monopoly for the product with the highest value $v_i q_i$ [2]. When $r > 1$, we show that the equilibria of the T-O market are given by monopolies for each product and an inner equilibrium (also characterized analytically). *We show that, when $r > 1$, the monopolies are stable, while the inner equilibrium is not.*

It is also useful to mention that different, theoretical and experimental, approaches to the use of social influence are present in the literature. For instance, Yuan and Hwarng [29] describe a demand-based pricing model under social influence and capture its behavior with a dynamical system that evolve to some stable or chaotic equilibria depending on the strength of the social signal. Stummer et al [23] introduces an agent-based model for repeat purchase decisions addressing different types of innovation diffusion and their perceived attributes; They also applied this methodology to an application concerned with second-generation biofuel.

3 The Trial-Offer Model

The paper builds on the work by Krumme et al [12] who propose a framework in which consumer choices are captured by a multinomial logit model whose product utilities depend on the product appeal, position bias, and a social influence signal representing past purchases. A marketplace consists of a set N of n items. Each item $i \in N$ is characterized by two values:

1. its *appeal* $a_i > 0$ which represents the inherent preference of trying item i;

2. its *quality* $q_i > 0$ which represents the conditional probability of purchasing item i given that it was tried.

This paper assumes that the appeals and the qualities are known. Abeliuk and al [2] have shown that these values can be recovered accurately and quickly, either before or during the market execution. The objective of the firm running this market is to maximize the total expected number of purchases. To achieve this, the key managerial decision of the firm is what is known as the ranking policy [2], which consists in deciding how to display the products in the market (e.g., where to display a product on a web page). Here we assume that, at the beginning of the market, the firm decides upon a ranking for the items, i.e., an assignment of items to positions in the marketplace. Each position j has a visibility v_j which represents the inherent probability of trying an item in position j. A ranking σ is a permutation of the items and $\sigma(i) = j$ means that item i is placed in position j ($j \in N$). When a customer enters the market, she observes all the items and their social signals based on the values of the previous purchases $d^t = (d_1^t, \ldots, d_n^t)$.

The vector ϕ^t of market shares at time t is computed in terms of the vector d^t, i.e.,

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

and

$$\phi^t \in \Delta^{n-1} = \{x = (x_1, \ldots, x_n) \in \Re^n \ \mid \ 0 \le x_i \le 1 \text{ and } \sum_{i=1}^n x_i = 1\}.$$

The consumer then selects an item to try. The probability that the customer tries item i is given by $P_i(\sigma, \phi^t)$ where

$$P_i(\sigma, \phi) = \frac{v_{\sigma(i)} f(\phi_i)}{\sum_{j=1}^n v_{\sigma(j)} f(\phi_j)} \tag{2}$$

and f is a continuous, positive, and nondecreasing function. This probability generalizes the multinomial logic model of Krumme et al [12]. Finally, the customer decides whether to buy the sampled item i and the probabilty that she purchases item i (after having tried it) is given by q_i. If item i is selected at time t, then the purchase vector becomes

$$d_j^{t+1} = \begin{cases} d_j^t + 1 & \text{if } j = i; \\ d_j^t & \text{otherwise.} \end{cases}$$

For simplicity, the vector d^0 is initialized with the product appeals, i.e., $d_i^0 = a_i$.[3] To analyze this process, we divide time into discrete periods such that each new period begins when a product has been purchased. Hence, the length of each time period is not constant.

In this paper, we are interested in characterizing how the market shares $\{\phi^t\}_{t>0}$ evolve over time for various functions f given a static ranking σ. We are particularly interested in study the

[3]The initialization can be justified by viewing the discrete dynamic process as an urn and balls process, where the appeals are the initial sets of balls.

asymptotic behavior of $\{\phi^t\}_{t>0}$ for the cases where $f(x) = x^r$ with $r > 0$ (This can be interpreted for instance, as displaying a social signal given by the square root of the number of past purchases, when $r = 0.5$). For notational simplicity, we assume that the ranking is fixed and is the identity function $\sigma(i) = i$ and omit it from the formulas. If the qualities and visibilities also satisfy $q_1 \geq \ldots \geq q_n$ and $v_1 \geq \ldots \geq v_n$, we obtain the quality ranking proposed in [26] but our results hold for any static ranking.

The following lemma relates the two phases of the T-O market and characterizes the probability that the next purchase is item i. It generalizes a similar result in [26].

Lemma 3.1. *If $f : \Delta^{n-1} \to \mathbb{R}$ is an continuous, positive, and nondecreasing function, then the probability $p_i(\phi)$ that the next purchase is the product i given the market share vector ϕ is given by*

$$p_i(\phi) \ = \frac{v_i q_i f(\phi_i)}{\sum_{j=1}^n v_j q_j f(\phi_j)}. \tag{3}$$

Proof. The probability that item i is purchased in the first step is given by

$$p_i^{1st}(\phi) = \frac{v_i f(\phi_i)}{\sum\limits_{j=1}^n v_j f(\phi_j)} q_i.$$

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

$$p_i^{2nd}(\phi) = \left(\frac{\sum\limits_{j=1}^n v_j f(\phi_j)(1-q_j)}{\sum\limits_{j=1}^n v_j f(\phi_j)} \right) \frac{v_i f(\phi_i)}{\sum\limits_{j=1}^n v_j f(\phi_j)} q_i.$$

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

$$p_i^{mth}(\phi) = \left(\frac{\sum\limits_{j=1}^n v_j f(\phi_j)(1-q_j)}{\sum\limits_{j=1}^n v_j f(\phi_j)} \right)^{m-1} \frac{v_i f(\phi_i)}{\sum\limits_{j=1}^n v_j f(\phi_j)} q_i. \tag{4}$$

Let $a = (\sum\limits_{j=1}^n v_j f(\phi_j) q_j)/(\sum\limits_{j=1}^n v_j f(\phi_j))$, Equation (4) becomes

$$p_i^{mth}(\phi) = \Big(1 - a\Big)^{m-1} \frac{v_i f(\phi_i)}{\sum\limits_{j=1}^n v_j f(\phi_j)} q_i.$$

Hence the probability that the next purchase is item i is given by

$$p_i(\phi) = \sum_{m=0}^{\infty} \Big(1 - a\Big)^{m} \frac{v_i f(\phi_i)}{\sum\limits_{j=1}^n v_j f(\phi_j)} q_i.$$

Since

$$\sum_{m=0}^{\infty} \left(1 - a\right)^{m} = \frac{1}{a},$$

the probability that the next purchase is item i is given by

$$p_i(\phi) = \frac{v_i q_i f(\phi_i)}{\sum_{j=1}^{n} v_j q_j f(\phi_j)}.$$

□

4 Trial-Offer Markets as Robbins-Monro Algorithms

This section establish some basic definitions and concepts that are useful in the rest of the paper. In particular, it shows that T-O markets can be modeled as Robbins-Monro algorithms and states some useful results. The results in this section are well-known in stochastic approximation.

Differential Equations: Consider the following differential equation

$$\frac{dy}{dt} = F(y) \tag{5}$$

where F is some vector field. The concept of equilibrium is central in the study of asymptotic behaviour for this type of equation.

Definition 4.1. *A vector $y^* \in \mathbb{R}^n$ is an equilibrium for differential equation (5) if $F(y^*) = 0$.*

We are interested in equilibria that satisfy (at least) certain stability criteria.

Definition 4.2. *An equilibrium y^* is said to be stable for Equation (5) if, given $\epsilon > 0$, there exists $\delta > 0$ such that $\|y(t) - y^*\| < \epsilon$ for all $t > 0$ and for all y such that $\|y - y^*\| < \delta$. We say that y^* is asymptotic stable if also satisfies*

$$\lim_{t\to\infty} y(t) = y^*.$$

We now show that the discrete stochastic process $\{\phi^t\}_{t\geq 0}$ can be modeled as a Robbins-Monro Algorithm (RMA) [9, 13].

Definition 4.3 (Robbins-Monro Algorithm)**.** *A Robbins-Monro Algorithm (RMA) is a discrete time stochastic processes $\{x^t\}_{t\geq 0}$ whose general structure is specified by*

$$x^{k+1} - x^k = \gamma^{k+1}[F(x^k) + U^{k+1}], \tag{6}$$

where

- *x^k takes its values in some Euclidean space (e.g., $\mathbb{R}^n$);*

- *γ^k is deterministic or stochastic and satisfies $\gamma^k > 0$, $\sum_{t\geq 1} \gamma^t = \infty$, and $\lim_{t\to\infty} \gamma^t = 0$ with probability 1;*
- *$F : \mathbb{R}^n \to \mathbb{R}^n$ is a deterministic continuous vector field;*
- *$\mathbb{E}[U^{k+1}|\mathcal{F}^k] = 0$, where $\mathcal{F}^k$ is the natural filtration of the entire process.*[4]

Recall that the probability that the next purchase is item i at time k is given by $p_i(\phi^k)$ and denote by e^k the random unit vector whose j^{th} entry is 1 if item j is the next purchase and 0 otherwise. The market share at time $k+1$ is given by

$$\phi^{k+1} = \frac{D^k \phi^k}{D^k + 1} + \frac{e^k}{D^k + 1}$$

where $D^k = \sum_{t=0}^{k} \sum_{i=1}^{n} d_i^t$. It follows that

$$\begin{aligned}
\phi^{k+1} &= \frac{(D^k + 1)\phi^k}{D^k + 1} - \frac{\phi^k}{D^k + 1} + \frac{e^k}{D^k + 1} \\
&= \phi^k + \frac{1}{D^k + 1}(e^k - \phi^k) \\
&= \phi^k + \frac{1}{D^k + 1}(\mathbb{E}[e^k|\phi^k] - \phi^k + e^k - \mathbb{E}[e^k|\phi^k]) \\
&= \phi^k + \frac{1}{D^k + 1}(p(\phi^k) - \phi^k + e^k - \mathbb{E}[e^k|\phi^k]).
\end{aligned}$$

This last equality can be reformulated as

$$\phi^{k+1} = \phi^k + \gamma^{k+1}(F(\phi^k) + U^{k+1}) \tag{7}$$

where $\gamma^{k+1} = \frac{1}{D^k+1}$, $F(\phi) = p(\phi) - \phi$, and $U^{k+1} = e^k - \mathbb{E}[e^k|\phi^k]$. Note that the function F captures the difference between the probabilities of purchasing the items (given the market shares) and the market shares at each time step. We can now prove that the discrete dynamic process $\{\phi^t\}_{t\geq 0}$ can be modeled as a Robbins-Monro algorithm.

Theorem 4.4. *The discrete stochastic dynamic process $\{\phi^t\}_{t\geq 0}$ can be modeled as the Robbins-Monro algorithm.*

Proof. The above derivation showed that $\{\phi^t\}_{t\geq 0}$ can be expressed through Equation (7), i.e.,

$$\phi^{k+1} = \phi^k + \gamma^{k+1}(F(\phi^k) + U^{k+1})$$

where $\gamma^{k+1} = \frac{1}{D^k+1}$, $F(\phi) = p(\phi) - \phi$, and $U^{k+1} = e_k - \mathbb{E}[e^k|\phi^k]$. It is easy to see that $\gamma^k > 0$, $\sum_{t\geq 1} \gamma^t = \infty$, $\lim_{t\to\infty} \gamma^t = 0$, and that $\mathbb{E}[U^{k+1}|\mathcal{F}^k]$ is equal to zero. □

[4] Because of this last condition on U^k, a Robbins-Monro Algorithm is also known as a *Martingale Difference Stochastic Approximation.*

Robbins-Monro algorithms are particularly interesting because, under certain conditions on x^t, γ^t, and U^t, their asymptotic behavior, i.e., the values of x^t when $t \to \infty$, is closely related to the asymptotic behavior of the following continuous dynamic process:

$$\frac{dx^t}{dt} = F(x^t), \quad x^t \in \mathbb{R}^n. \tag{8}$$

This idea, called *the ODE Method*, was introduced by [16] and has been extensively studied (e.g., [5, 9, 13]). Consider again the RMA $\{x^t\}_{t\geq 0}$ defined in (6) and the following hypotheses:

H1: $\sup_t \mathbb{E}[\|U^{t+1}\|^2] < \infty$;

H2: $\sum_t (\gamma^t)^2 < \infty$;

H3: $\sup_t \|x^t\| < \infty$.

We will now present a theorem establishing the connection between the discrete stochastic process (6), and the continuous process defined (8) but we first need a key concept.

Definition 4.5. *A closed set A is an Internally Chain Transitivity (ICT) if for every $\epsilon > 0$ and every $x, y \in A$, there is a set $\{x_i\}_{i=0}^n \subset A$, such that $x_0 = x, x_n = y$, and $\|F(x_i) - x_{i+1}\| < \epsilon$, $\forall i \in \{0, ..., n-1\}$.*

The following theorem links the behaviour of the limit set $L\{x_t\}_{t\geq 0}$ of any sample path $\{x_t\}_{t\geq 0}$ for Equation (6) and the limit sets of the solution to Equation (8). The result is due to [4].

Theorem 4.6. *Let $\{x_t\}_{t\geq 0}$ be a Robbins-Monro algorithm (6) satisfying hypotheses $H1 - H3$ where F is locally Lipschitz (e.g., a $\mathcal{C}^1$ function). Then, with probability 1, the limit set $L\{x_t\}_{t\geq 0}$ is internally chain transitive for Equation (8).*

Remark 4.7. ICT sets include equilibria, periodic orbits of (8), but possibly more complicated sets. The noise U^{k+1} of (6) may push the stochastic process away from unstable sets. We also notice that Theorem 4.6 is valid for very general functions $F(x) = p(x) - x$, as the only requirement is to be locally Lipschitz.

Since hypotheses H1, H2, and H3 are all satisfied by the discrete stochastic dynamic process $\{\phi^t\}_{t\geq 0}$, it remains to study the structure of the ICT set of Equation (8). Starting from here we will focus in the case that the social signal $f(x)$ is given by $f(x) = x^r$, with $r > 0$.

5 Equilibria of Trial-Offer Markets

This section characterizes the equilibria and the asymptotic behaviour of continuous dynamics

$$\frac{d\phi^t}{dt} = p(\phi^t) - \phi^t \quad (\phi^t \in \Delta^{n-1}) \tag{9}$$

associated with the RMA (7). For simplicity, we remove the visibilities by stating $\bar{q}_j = v_j q_j$. We are mainly interested in the case where $f(x) = x^r$ with $(0 < r < 1)$, since the case $r = 1$ has been settled in earlier work.

Theorem 5.1. *Let $f(x) = x^r, r > 0, r \neq 1$. Then, there is a unique equilibrium to Equation (9) in the interior $int(\Delta^{n-1})$ of the simplex Δ^{n-1} specified by*

$$\phi^* = \frac{1}{\sum_j \overline{q}_j^{\frac{1}{1-r}}}[\overline{q}_1^{\frac{1}{1-r}}, \cdots, \overline{q}_n^{\frac{1}{1-r}}]$$

The remaining equilibria are on the boundary of the simplex.

Proof. We consider the set Q defined as $Q = \{i \in \{1, ..., n\} \mid \phi_i = 0\}$. An equilibrium to (9) must satisfy $p_i(\phi) = \phi_i$, i.e.,

$$\frac{\overline{q}_i(\phi_i)^r}{\sum_{j=0}^n \overline{q}_j(\phi_j)^r} = \phi_i.$$

For $i \notin Q$, we have

$$\overline{q}_i(\phi_i)^{r-1} = \sum_{j \notin Q} \overline{q}_j(\phi_j)^r$$

and, for all $i, k \notin Q$, we also have

$$\overline{q}_i(\phi_i)^{r-1} = \sum_{j \notin Q} \overline{q}_j(\phi_j)^r = \overline{q}_k(\phi_i)^{r-1}$$

which is equivalent to

$$\overline{q}_i(\phi_i)^{r-1} = \overline{q}_k(\phi_j)^{r-1} \Leftrightarrow \phi_i = \left(\frac{\overline{q}_k}{\overline{q}_i}\right)^{\frac{1}{r-1}} \phi_k. \tag{10}$$

By summing for all $i \notin Q$, we obtain

$$1 = \sum_{i \notin Q} \phi_i = \frac{\phi_k}{\overline{q}_k^{1/(1-r)}} \sum_{i \notin Q} \overline{q}_i^{1/(1-r)}$$

and hence

$$\phi_k = \frac{\overline{q}_k^{1/(1-r)}}{\sum_{i \notin Q} \overline{q}_i^{1/(1-r)}}.$$

We notice that, when $Q = \emptyset$, the equilibrium lives in the interior of the simplex $int(\Delta^{n-1})$. On the other hand, if $|Q| = n - 1$, then the equilibrium is one of the vertices of the simplex. □

Observe that the equilibrium $\phi^* \in int(\Delta^K)$ for the case $0 < r < 1$ has some very interesting properties: It is unique, which means that the market is completely predictable. Moreover, if $\overline{q}_i \geq \overline{q}_j$, then $\phi_i^* \geq \phi_j^*$, which endows the market with a basic notion of *fairness*. Finally, the market is not a monopoly: All the market shares are strictly positive for the equilibrium ϕ^*.

Our next result characterizes the ICT of the continuous dynamics. We start with a useful lemma which indicates that submarkets can also be modeled as RMAs.

Lemma 5.2. *Consider a T-O market defined by n items and the submarket obtained by considering only $n - 1$ items. Then this submarket can also be modeled by an RMA.*

Proof. Let $\Phi^t = [\phi_1^t, \phi_2^t, \cdots, \phi_n^t]$ be the market share for the n-item T-O market at stage t. Consider a new process $\{\Psi^t\}_{t\geq 0}$ consisting of $n-1$ products only. We show that this process can also be modeled as a RMA. The key is to prove that the probability of purchasing product j in stage t follows Equation (3). Consider any item $i \in \{1, .., n\}$ such that $\phi_i^t \neq 1$. Without loss of generality, assume that $i = n$, define

$$\psi_i^t = \frac{\phi_i^t}{1-\phi_n^t}, \quad (i < n),$$

and consider the following events:

- $A = \{\text{product } n \text{ is not purchased at stage } t\}$
- $B = \{\text{product } j \neq n \text{ is purchased at stage } t\}$.

Since $B \subseteq A$, $Pr[B \cap A] = Pr[B] = \frac{\overline{q}_j(\phi_j^t)^r}{\sum_{i=1}^n \overline{q}_i(\phi_i^t)^r}$. On the other hand

$$Pr[A] = 1 - \frac{\overline{q}_n(\phi_n^t)^r}{\sum_{i=1}^n \overline{q}_i(\phi_i^t)^r} = \frac{\sum_{j=1}^{n-1} \overline{q}_j(\phi_j^t)^r}{\sum_{i=1}^n \overline{q}_i(\phi_i^t)^r},$$

and therefore

$$Pr[B|A] = \frac{Pr[B \cap A]}{Pr[A]} = \frac{\overline{q}_j(\phi_j^t)^r}{\sum_{i=1}^{n-1} \overline{q}_i(\phi_i^t)^r} \cdot \frac{(1-\phi_n^t)^r}{(1-\phi_n^t)^r} = \frac{\overline{q}_j(\psi_j^t)^r}{\sum_{i=1}^{n-1} \overline{q}_i(\psi_i^t)^r}.$$

Since $\psi_i^t \geq 0$ and $\sum_{i=1}^{n-1} \psi_i^t = \sum_{i=1}^{n-1} \frac{\phi_i^t}{1-\phi_n^t} = \frac{1}{1-\phi_n^t} \sum_{i=1}^{n-1} \phi_i^t = 1$, the ψ_i^t are well-defined market shares. Since all the remaining properties from the original model Φ^t still hold, $\{\Psi^t\}_{t\geq 0}$ can be modeled as a $n-1$ dimensional RMA. □

We are now in position to prove the main result of this paper.

Theorem 5.3. *Under the social signal $f(x) = x^r, 0 < r < 1$ with $\phi^0 \in int(\Delta^n)$, the RMA $\{\phi^t\}_{t>0}$ converges to ϕ^* almost surely.*

Proof. The proof studies the asymptotic behaviour of the solutions of the following ODE:

$$\frac{d\phi^t}{dt} = p(\phi^t) - \phi^t. \tag{11}$$

Equation (11) is equivalent to

$$\frac{d\phi_i^t}{dt} = \frac{\overline{q}_i(\phi_i^t)^r}{\sum_j \overline{q}_j(\phi_j^t)^r} - \phi_i^t, \quad i \in \{1, \cdots, n\}.$$

Hence, we have

$$
\begin{aligned}
\frac{\overline{q}_i(\phi_i^t)^r}{\sum_j \overline{q}_j(\phi_j^t)^r} &= \frac{d\phi_i^t}{dt} + \phi_i^t, \\
\frac{1}{\sum_j \overline{q}_j(\phi_j^t)^r} &= \frac{1}{\overline{q}_i(\phi_i^t)^r}[\frac{d\phi_i^t}{dt} + \phi_i^t] \quad \text{if } \phi_i^t \neq 0, \\
\frac{1}{\overline{q}_i(\phi_i^t)^r}[\frac{d\phi_i^t}{dt} + \phi_i^t] &= \frac{1}{\overline{q}_j(\phi_j^t)^r}[\frac{d\phi_j^t}{dt} + \phi_j^t] \quad \forall i, j \in \{1, \cdots, n\}, \\
\overline{q}_i^{-1}(\phi_i^t)^{-r}[\frac{d\phi_i^t}{dt} + \phi_i^t] &= \overline{q}_j^{-1}(\phi_j^t)^{-r}[\frac{d\phi_j^t}{dt} + \phi_j^t], \\
\overline{q}_i^{-1}[(\phi_i^t)^{-r}\frac{d\phi_i^t}{dt} + (\phi_i^t)^{1-r}] &= \overline{q}_j^{-1}[(\phi_j^t)^{-r}\frac{d\phi_j^t}{dt} + (\phi_j^t)^{1-r}], \\
e^{(1-r)t}(1-r)\overline{q}_i^{-1}[(\phi_i^t)^{-r}\frac{d\phi_i^t}{dt} + (\phi_i^t)^{1-r}] &= e^{(1-r)t}(1-r)\overline{q}_j^{-1}[(\phi_j^t)^{-r}\frac{d\phi_j^t}{dt} + (\phi_j^t)^{1-r}], \\
\frac{d}{dt}\left[e^{(1-r)t}\overline{q}_i^{-1}(\phi_i^t)^{1-r}\right] &= \frac{d}{dt}\left[e^{(1-r)t}\overline{q}_j^{-1}(\phi_j^t)^{1-r}\right]
\end{aligned}
$$

where the fourth equivalence is obtained by multiplying both sides with $\mu(t) = (1-r)e^{(1-r)t}$. Taking the integral $\int_0^t dt$ of the last expression gives

$$
e^{(1-r)t}\overline{q}_i^{-1}(\phi_i^t)^{1-r} - \overline{q}_i^{-1}(\phi_i^0)^{1-r} = e^{(1-r)t}\overline{q}_j^{-1}(\phi_j^t)^{1-r} - \overline{q}_j^{-1}(\phi_j^0)^{1-r} \tag{12}
$$

and hence

$$
\frac{(\phi_i^t)^{1-r}}{\overline{q}_i} - \frac{(\phi_j^t)^{1-r}}{\overline{q}_j} = e^{(r-1)t}\left[\frac{(\phi_i^0)^{1-r}}{\overline{q}_i} - \frac{(\phi_j^0)^{1-r}}{\overline{q}_j}\right]. \tag{13}
$$

Consider Equation (13):

- if, for some $i \neq j$, $\frac{(\phi_i^0)^{1-r}}{\overline{q}_i} = \frac{(\phi_j^0)^{1-r}}{\overline{q}_j}$, then $\frac{(\phi_i^t)^{1-r}}{\overline{q}_i} = \frac{(\phi_j^t)^{1-r}}{\overline{q}_j}$, for all t;
- if $\frac{(\phi_i^0)^{1-r}}{\overline{q}_i} \neq \frac{(\phi_j^0)^{1-r}}{\overline{q}_j}$, the right-hand side of Equation (13) goes to zero as $t \to \infty$ (because $r < 1$) and hence
$$
\lim_{t\to\infty} \frac{(\phi_i^t)^{1-r}}{\overline{q}_i} - \frac{(\phi_j^t)^{1-r}}{\overline{q}_j} = 0. \tag{14}
$$

We now prove by induction that the limits for the market shares exist. Consider first the case of 2 products. Since $\phi_2^t = (1-\phi_1^t)$, the market can be modeled as a one-dimensional RMA. By Theorem 1 in [19], the RMA converges under mild conditions trivially satisfied here. Assume now that a RMA with $k-1$ products converges and consider a market with k products. By Lemma 5.2, given a k-dimensional RMA $\Phi^t = [\phi_1^t, \phi_2^t, \cdots, \phi_k^t]$ we can create a $k-1$ dimensional RMA $\{\Psi^t\}_{t\geq 0}$ given by $\psi_i^t = \frac{\phi_i^t}{1-\phi_k^t}, \quad i < k,$. By induction hypothesis, $\psi_i = \lim_{t\to\infty} \psi_i^t$ exists for all $i < k$ and therefore Equation (13) is equivalent to

$$
\frac{(\phi_k^t)^{1-r}}{\overline{q}_k(1-\phi_k^t)^{1-r}} - \frac{(\psi_i^t)^{1-r}}{\overline{q}_i} = \frac{e^{(r-1)t}}{(1-\phi_k^t)^{1-r}}\left[\frac{(\phi_k^0)^{1-r}}{\overline{q}_k} - \frac{(\phi_i^0)^{1-r}}{\overline{q}_i}\right]. \tag{15}
$$

The right hand side of (15) goes to 0 when $t \to \infty$ and $\lim_{t\to\infty} \frac{(\psi_i^t)^{1-r}}{\overline{q}_i}$ exists. Hence $\lim_{t\to\infty} \frac{(\phi_n^t)^{1-r}}{\overline{q}_n(1-\phi_n^t)^{1-r}}$ also exists.

Now denote by ϕ_j the limit of ϕ_j^t for all $j \in \{1, \cdots, n\}$. Using Equation (14), the following equation holds for all $i, j \in \{1, \cdots, n\}$:

$$\frac{\phi_i^{1-r}}{\overline{q}_i} = \frac{\phi_j^{1-r}}{\overline{q}_j}. \tag{16}$$

Observe that, if there exists $l \in \{1, \cdots, n\}$ such that $\phi_l = 0$, Equation (16) implies that $\phi_i = 0$ for all i which is impossible since they sum to 1. Hence the limit process has strictly positive components and Equation (16) is equivalent to

$$\phi_i = \frac{\phi_j}{\overline{q}_j^{1/(1-r)}} \overline{q}_i^{1/(1-r)} \tag{17}$$

which is the equation that defines ϕ^* in Theorem 5.1 (see Equation (10)). As a result, when $\Phi_0 \in int(\Delta^{n-1})$, the only ICT set for the ODE (11) is the equilibrium ϕ^* and the RMA converges to ϕ^*. □

Consider now the case $r > 1$. We know that Theorem 5.1 is valid to characterize the equilibria. However, the dynamic behaviour is completely different due to the strength of the social signal. In particular, the inner equilibrium Φ^* is now unstable. The proof uses the following characterization of unstability for one-dimensional RMA due to Renlund.

Proposition 5.4. *([19]) Consider a one-dimentional RMA with $F(x) = p(x) - x$. A point x^* is unstable if there exists a neighbourhood N_{x^*} around x^* such that $F(x)[x - x^*] \geq 0$ for all $x \in N_{x^*}$.*

Theorem 5.5. *For a social signal $f(x) = x^r$ with $r > 1$, the inner equilibrium is unstable.*

Proof. Let $x = [x_1, x_2]$ be the market share for the case of 2 products (at any time t). Since $x_2 = 1 - x_1$, $x_i^* = \frac{q_i^{1/(1-r)}}{q_1^{1/(1-r)} + q_2^{1/(1-r)}}$. Let $c_i = q_i^{1/(1-r)}$ ($\Rightarrow q_i = c_i^{1-r}$). We have

$$\begin{aligned}
[p(x_1) - x_1][x_1 - x_1^*] &= \left[\frac{c_1^{1-r}x_1^r}{c_1^{1-r}x_1^r + c_2^{1-r}(1-x_1)^r} - x_1\right]\left[x_1 - \frac{c_1}{c_1+c_2}\right] \\
&= \frac{\left[c_1^{1-r}x_1^r(1-x_1) - c_2^{1-r}(1-x_1)^r x_1\right][x_1(c_1+c_2) - c_1]}{(c1+c2)(c_1^{1-r}x_1^r + c_2^{1-r}(1-x_1)^r)} \\
&= \frac{[x_1(1-x_1)][c_1^{1-r}x_1^{r-1} - c_2^{1-r}(1-x_1)^{r-1}][x_1(c_1+c_2) - c_1]}{(c_1+c_2)(c_1^{1-r}x_1^r + c_2^{1-r}(1-x_1)^r)}.
\end{aligned}$$

Since the denominator is always positive and $x_1(1-x_1) \geq 0$, we only need to study the sign of $[c_1^{1-r}x_1^{r-1} - c_2^{1-r}(1-x_1)^{r-1}][x_1(c_1+c_2) - c_1]$. Then,

$$\begin{aligned} x_1(c_1+c_2) - c_1 > 0 &\Leftrightarrow x_1 > \frac{c_1}{c_1+c_2} \\ &\Leftrightarrow 1 - x_1 < 1 - \frac{c_1}{c_1+c_2} = \frac{c_2}{c_1+c_2} \\ &\Leftrightarrow [c_1^{1-r}x_1^{r-1} - c_2^{1-r}(1-x_1)^r x - 1] > [c_1^{1-r}\left(\frac{c_1}{c_1+c_2}\right)^{r-1} - c_2^{1-r}\left(\frac{c_2}{c_1+c_2}\right)^{r-1}] \end{aligned}$$

Since $c_1^{1-r}\left(\frac{c_1}{c_1+c_2}\right)^{r-1} - c_2^{1-r}\left(\frac{c_2}{c_1+c_2}\right)^{r-1} = \frac{1}{(c_1+c_2)^{r-1}} - \frac{1}{(c_1+c_2)^{r-1}} = 0$, $[p(x_1) - x_1][x_1 - x_1^*] \geq 0$ for all $x_1 \in (0,1)$. And hence x^* is an unstable equilibrium.

This proof can be easily extended for the general case of n products, just applying the previous argument recursively to $x_1 = i, x_2 = \{j \in \{1, \cdots, n\} : j \neq i\}$. □

Theorem 5.6. *Consider the social signal $f(x) = x^r$ with $r > 1$. The RMA $\{\phi^t\}_{t\geq 0}$ converges almost surely to one of the equilibria $\phi \in Z_F := \{x \in \Delta^{n-1} : p(x) - x = 0\}$.*

Proof. The analysis of the ODE is the same as in Theorem 5.3 since the only restriction in the proof is $r \neq 1$. However, the interpretation of Equation (13) changes when $r > 1$.

We let $H_{i,0} := \dfrac{(\phi_i^0)^{1-r}}{\overline{q}_i}$ and order the products in decreasing order of $H_{i,0}$. Let $h : \{1,..,n\} \to \{1,..,n\}$ be the permutation that defines this order and denotes by h^{-1} its inverse function, i.e., $h^{-1}(i) = j$ means that product j is in the i-th position in permutation h. We have that $H_{h^{-1}(1),0} \geq \cdots \geq H_{h^{-1}(n),0}$. Define the following sets:

- $Q_0 = \{i \in \{1,..,n-1\} : H_{h^{-1}(i),0} = H_{h^{-1}(i+1),0}\}$,
- $Q_1 = \{i \in \{1,..,n-1\} : H_{h^{-1}(i),0} > H_{h^{-1}(i+1),0}\}$,

and consider the following case analysis:

i) If $|Q_0| = n-1$, then $H_{h^{-1}(i),0} = H_{h^{-1}(i+1),0}$ for all $1 < i < n-1$. By Equation (13), $\dfrac{(\phi^t_{h^{-1}(i)})^{1-r}}{\overline{q}_{h^{-1}(i)}} = \dfrac{(\phi^t_{h^{-1}(i+1)})^{1-r}}{\overline{q}_{h^{-1}(i+1)}}$, for all $t > 0$ and for all $1 < i < n-1$, which is again, the inner equilibrium.

ii) If $0 < |Q_0| < n-1$, select $i \notin Q_0$. Equation (13) implies that

$$\lim_{t\to\infty} \frac{(\phi^t_{h^{-1}(i)})^{1-r}}{\overline{q}_{h^{-1}(i)}} - \frac{(\phi^t_{h^{-1}(i+1)})^{1-r}}{\overline{q}_{h^{-1}(i+1)}} = \infty,$$

because $r > 1$ and hence $e^{(r-1)t} \to \infty$ when $t \to \infty$. It follows that $\lim_{t\to\infty} \phi^t_{h^{-1}(i)} = 0$ and the RMA necessarily converges to one of the equilibria that live in the boundary of the simplex (see Theorem 5.1).

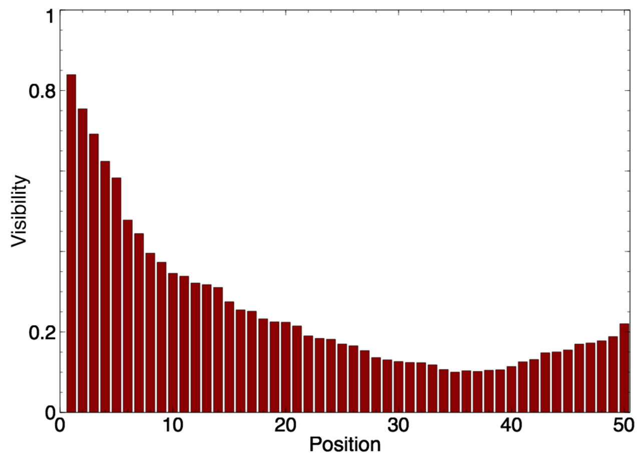

Figure 1: The visibility v_p (y-axis) of position p in the song list (x-axis) where $p = 1$ is the top position and $p = 50$ is the bottom position of the list which is displayed in a single column.

iii) If $|Q_0| = 0$ then $|Q_1| = n - 1$, Using a similar reasoning as in case (ii), it follows that $\lim_{t\to\infty} \phi^t_{h^{-1}(i)} = 0$ for all $1 < i < n - 1$ and, since $\phi^t \in \Delta^{n-1}$ for all t, $\lim_{t\to\infty} \phi^t_{h^{-1}(n)} = 1$.

As a result, the only ICT for the differential equation (11) are equilibria and the RMA $\{\phi^t\}_{t\geq 0}$ converges almost surely to one of them. □

It is important to observe that, in the case $r > 1$, the initial condition affects the entire dynamics. This is in contrast with the case $r < 1$ for which the long-term behaviour is independent of the initial condition. This has obviously fundamental consequences for the predictability of the market. Note also that, by Theorem 5.5, ϕ^* is an unstable equilibrium when $r > 1$ and hence the RMA $\{\phi^t\}_{t\geq 0}$ most likely converges to one of the boundary equilibria $\phi \in Z_F \setminus \{\phi^*\}$.

6 Agent-Based Simulation Results

We now report results from an agent-based simulation to highlight and complement the theoretical analysis. The agent-based simulation uses settings that model the MusicLab experiments discussed in [2, 12, 20]. 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 [12] uses the consumer choice preferences described in Section 3.

The Simulation Setting The agent-based simulation aims at emulating 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 $p_i(\sigma, \phi)$, where σ is the ranking proposed by the policy under evaluation and ϕ represents the market shares;

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 number of downloads of song i, i.e., $d_i^{t+1} = d_i^t + 1$, changing the market shares. Otherwise, $d_i^{t+1} = d_i^t$.

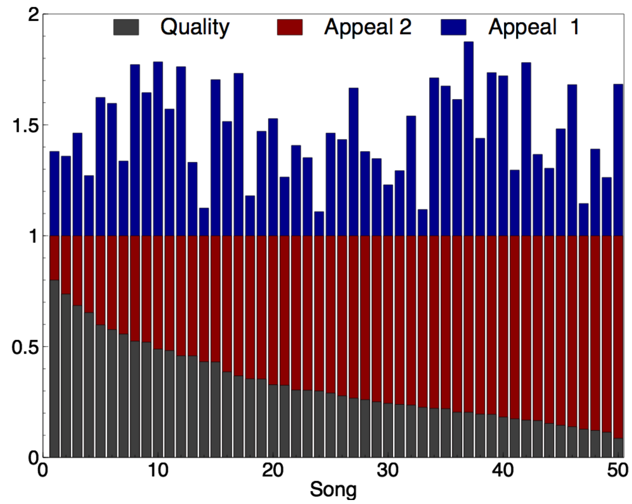

Figure 2: The Quality q_i (gray) and Appeal a_i (red and blue) of song i in the two settings. The settings only differ in the appeal of songs, and not in the quality of songs. In the first setting, the quality and the appeal for the songs were chosen independently according to a Gaussian distribution normalized to fit between 0 and 1. The second setting explores an extreme case where the appeal is negatively correlated with the quality used in setting 1. In this second setting, the appeal and quality of each song sum to 1.

Every $t > 0$ iterations, a new list σ may be recomputed if the ranking policy is dynamic (e.g., the popularity ranking). In this paper, the simulation setting focuses mostly on two policies for ranking the songs:

- The *quality ranking* (Q-rank) that assigns the songs in decreasing order of quality to the positions in decreasing order of visibility (i.e., the highest quality song is assigned to the position with the highest visibility and so on);
- The *popularity ranking* (D-rank) that assigns the songs in decreasing order of popularity (i.e., d_i^t) to the positions in decreasing order of visibility (i.e., the most popular song is assigned to the position with the highest visibility and so on);

Note that the popularity ranking was used in the original MusicLab, while the quality ranking is a static policy: the ranking remains the same for the entire simulation. The simulation setting, which aims at being close to the MusicLab experiments, considers 50 songs and simulations with 20,000 steps unless stated otherwise. The songs are displayed in a single column. The analysis in [12] indicated 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. Figure 1 depicts the visibility profile based on these guidelines, which is used in all computational experiments. The paper also uses two settings for the quality and appeal of each song, 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. Unless stated otherwise, the results were obtained by averaging the results of $W = 400$ simulations.

6.1 Convergence

We first illustrate the convergence of the market for various popularity signals ($r < 1$) using the quality ranking. In order to visualize the results, we focus on only 5 songs, where the qualities, appeals, and visibilities are given by

q = [0.8, 0.72, 0.68, 0.65, 0.60]
a = [0.38, 0.35, 0.46, 0.27, 0.62]
v = [0.8, 0.75, 0.69, 0.62, 0.58].

The simulation is run for 10^5 time steps for the social signals $f(x) = x^r (r \in \{0.1, 0.25, 0.5, 0.75\})$ and Figure 3 depicts the simulation results. Observe that the equilibrium ϕ^* (dashed lines) changes because it depends of the value of r. Interestingly, for social signals with $r \leq 0.5$, the convergence of the process seems to occur around 10^4 time steps even when they start with a strong initial distortion due to the appeals of the songs. The simulations show clear differences in behavior depending on r and, when r moves closer to 1, the market tends to exhibit a monopolistic behaviour for the song with the best quality (confirming the results obtained in [26]).

Figure 4 shows how the market is distributed in the equilibrium for 6 songs also taken from the MusicLab. The qualities, appeals, and visibilities are given by

q = [0.8,0.72,0.65,0.57,0.52,0.4887]
a = [0.38,0.36,0.27,0.60,0.77,0.78]
v = [0.8,0.75,0.62,0.48,0.40,0.35]

In the figure, the songs are ordered from left to right by increasing quality and the social signals are of the form $f(x) = x^r$ ($r \in \{0.1, 0.25, 0.5, 0.75\}$). Songs with better qualities (i.e., the top 2 songs on the right in this case) have larger market shares and their market shares increase with r. In contrast, the market shares of the lower-quality songs (i.e., the bottom four songs on the left in this case) decrease when r increases. These results indicate that social influence has a beneficial effect on the market: It drives customers towards the better products, while not going to a monopoly as long as $r < 1$.

6.2 Market Predictability

This section depicts the predictability of the market for various values of r and the number of downloads per song as a function of its quality. Figures 5 and 6 depict the results for the two quality/appeal settings discussed previously. The figures display the results of 200 experiments for each setting. Each simulation contributes 50 data points, i.e., the number of downloads for each song and all the data points for the 200 simulations are displayed in the figures. In the plots, the x-axis represents the song qualities and the y-axis the number of downloads. A dot at location (q, d) indicates that the song with quality q had d downloads in a simulation. Obviously, there can be several dots at the same location. For $r \in \{0.5, 0.75\}$, the market is highly predictable and there is little variation in the song downloads. For $r = 1$, the market converges to a monopoly for the song of highest quality, confirming the results from [2, 26]. Finally, for $r = 2$, the market exhibits

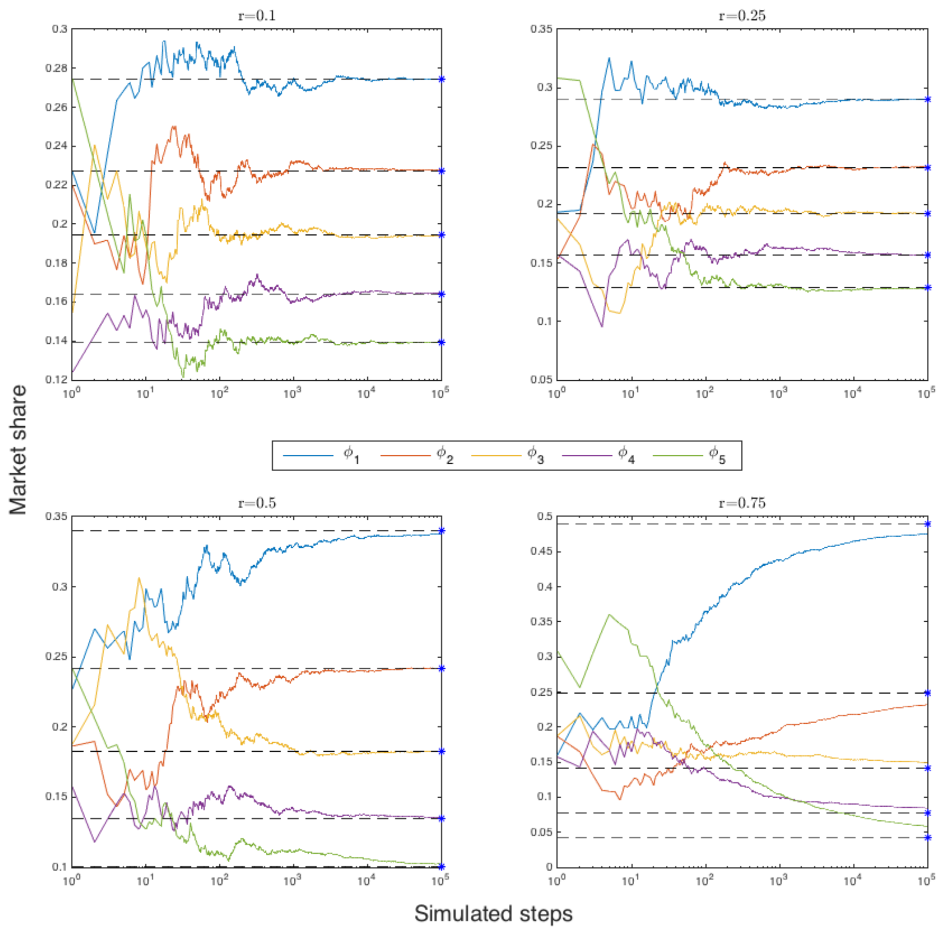

Figure 3: Evolution of market shares of 5 songs using a social signal $f(x) = x^r$, $r \in \{0.1, 0.25, 0.5, 0.75\}$. Dashed lines are the values of the equilibrium for each song.

significant unpredictability, as suggested by the theoretical results. In this case, the equilibria are monopolies for various songs but it is hard to predict which song will dominate the market. Note also that the unpredictability of the market increases significantly for $r = 2$ when the appeal and quality of the songs are negatively correlated. This is not the case for $r \in \{0.5, 0.75\}$ and less so for $r = 1$.

Figure 7 compares the predictability of Q-rank and D-rank for the first setting of Quality/Appeal. For each ranking, two different social signals were used ($r = 0.5$ and $r = 1$) and the figure displays the result of 50 experiments, consisting in 1 million of simulated steps. Two phenomena can be observed. First, sublinear signals seem to help the D-rank, making the outcome less chaotic (first column). Second, Q-rank clearly performs better than D-rank and exhibits much less unpredictability.

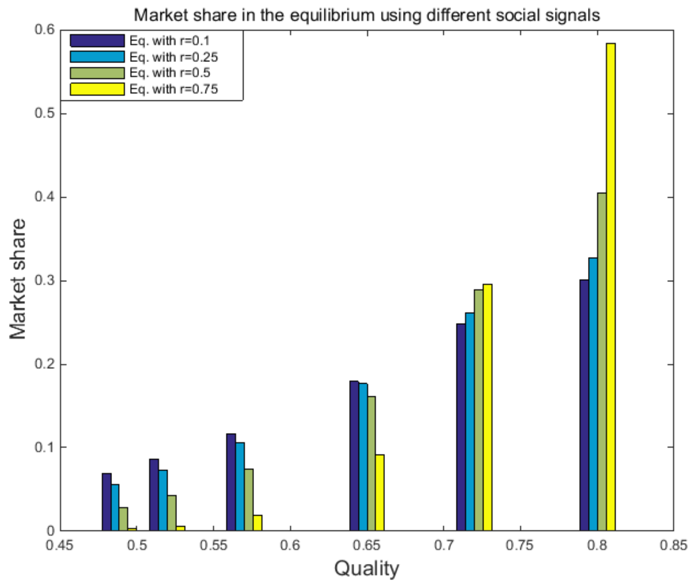

Figure 4: Market share of 6 songs respect their qualities, using a social signal $f(x) = x^r, \quad r \in \{0.1, 0.25, 0.5, 0.75\}$.

6.3 Performance of the Market

Figures 8 and 9 report results about the performance of the markets as a function of the social influence signals. The figures report the average number of downloads over time for the quality and popularity rankings as a function of the social signals. There are a few observations that deserve mention.

1. For the quality ranking, the expected number of downloads increases with the strength of the social signal as r converges to 1. The equilibrium when $r = 1$ is optimal asymptotically and assigns the entire market share to the song of highest quality. When $r = 2$, the situation is more complicated. In the second setting, when the simulation is run for more time steps, the market efficiency decreases slightly compared to $r = 1$, which is consistent with the theory since there is no guarantee that the monopoly for $r > 1$ is for the song of highest quality.

2. For the quality ranking, the improvements in market efficiency from social influence are quite substantial.

3. The popularity ranking is always dominated by the quality ranking and the benefits of the quality ranking increase as r increases.

4. The popularity ranking behaves in a catastrophic way when $r = 2$ in the second setting.

7 Additional Properties of Sublinear Social Signals

This section makes three additional observations about sublinear social signals.

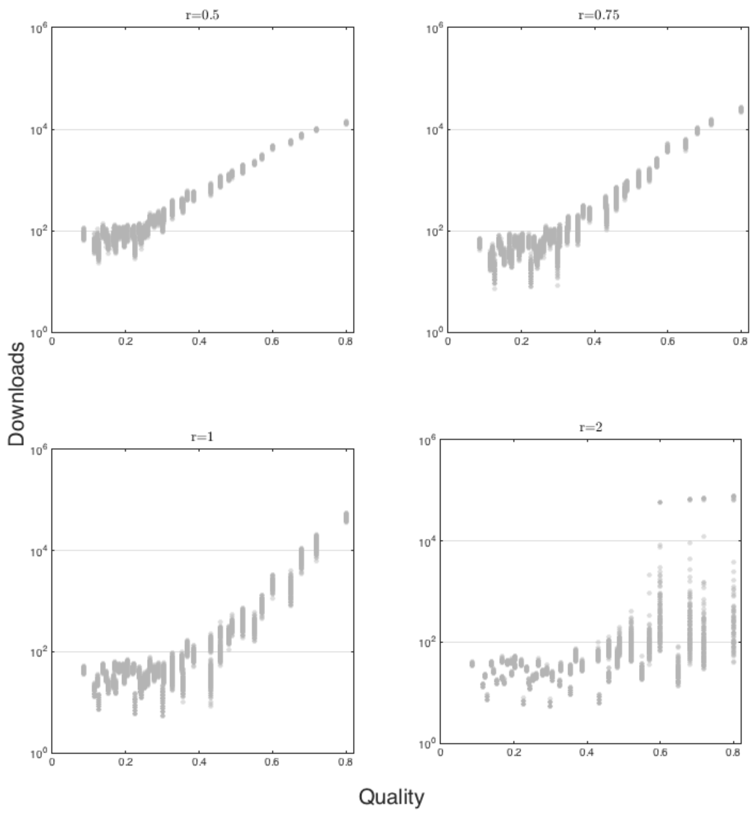

Figure 5: Distribution of Downloads Versus the Qualities, using Social Signals $f(x) = x^r$, $r \in \{0.5, 0.75, 1, 2\}$. The results are for the first setting where the quality and appeal of each song are not correlated. The songs are ordered by increasing quality along the x-axis. The y-axis is the number of downloads.

The Benefits of Social Influence A linear social signal always benefits the market in expectation when the performance ranking, the quality ranking, or any other static ranking is used. Unfortunately, sublinear social signals are not always beneficial to the market contrary to the case of linear social signal $r = 1$ [2, 26]. Consider, once again, the quality ranking and assume that $q_1 \geq \ldots \geq q_n$. When there is no social signal, the probability of trying product i is given by

$$P_i^- = \frac{v_i a_i}{\sum_{j=1}^n v_j a_j}$$

if quality ranking is used and the expected number of purchases is

$$\sum_{i=1}^n P_i^- q_i.$$

On the other hand, with a social signal, the probability of trying product i at time t is

$$P_i(\phi^t) = \frac{v_i f(\phi_i^t)}{\sum_{j=1}^n v_j f(\phi_j^t)}$$

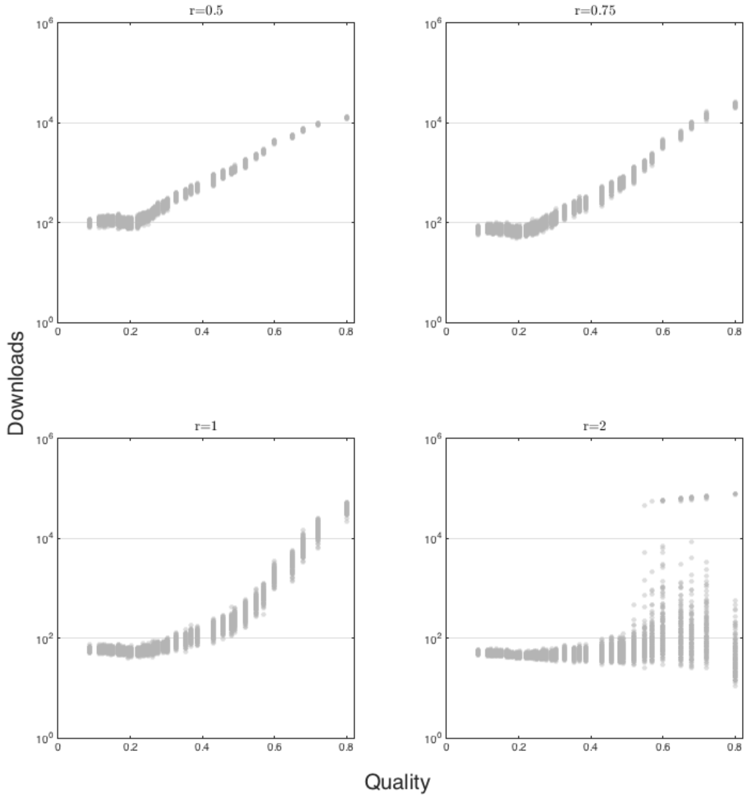

Figure 6: Distribution of Downloads Versus the Qualities, using Social Signals $f(x) = x^r$, $r \in \{0.5, 0.75, 1, 2\}$. The results are for the first setting where the quality and appeal of each song are negatively correlated. The songs are ordered by increasing quality along the x-axis. The y-axis is the number of downloads.

.

and the expected number of purchases at the equilibrium is given by

$$\sum_{i=1}^{n} P_i(\phi^*) q_i = \sum_{i=1}^{n} \frac{v_i q_i f(\phi_i^*)}{\sum_{j=1}^{n} v_j f(\phi_j^*)}.$$

Example 7.1. *Consider a 2-dimensional T-O market with social signal $f(x) = x^{0.5}$, where the qualities, visibilities, and appeals are given by*

- $q_1 = 1$, $q_2 = 0.4$,
- $v_1 = 1$, $v_2 = 0.8$,
- $a_1 = 1$, $a_2 = 0.1$.

The expected number of purchases with the social signal at the equilibrium is about 0.88 with social influence and about 0.95 without the social signal.

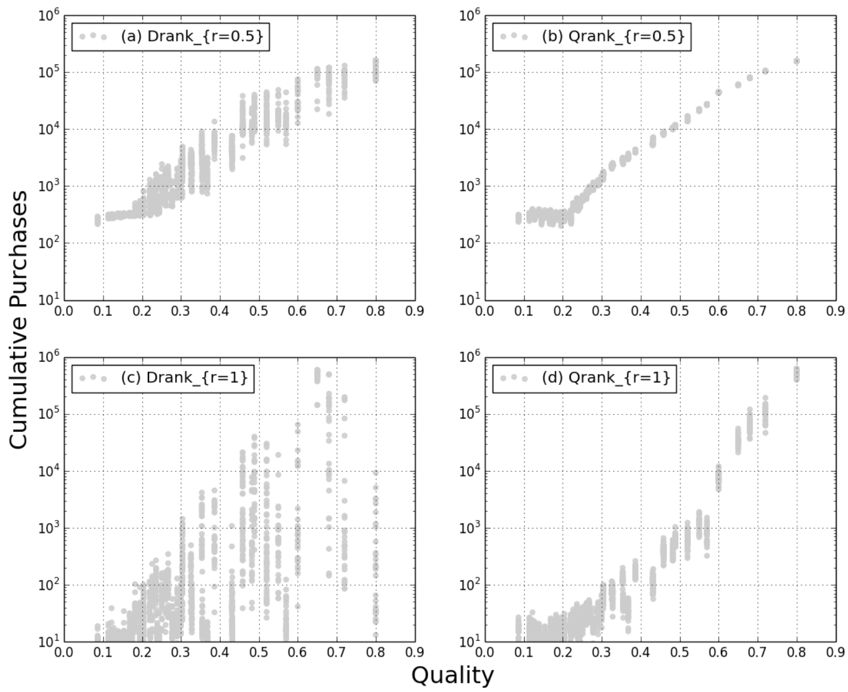

Figure 7: Distribution of purchases versus product qualities for 50 experiments with 1 million users. Figures (a) and (b) use a social signal $f(x) = x^{0.5}$, Figure (a) shows the results for the popularity ranking and Figure (b) for the quality ranking. Figures (c) and (d) use the social signal $f(x) = x$, Figure (c) shows the results for the popularity ranking and Figure (d) for the quality ranking.

The simple example is illuminating. It shows that the value of a sublinear social signal lies primarily in its ability to correct a misalignment of appeal and quality. If qualities and appeals are positively correlated, there is little use for a sublinear social influence signal. In contrast, when $r = 1$, social influence drives the market towards a monopoly, which leads to an asymptotically optimal market that assigns the entire market share to the highest quality product (which may be undesirable in practice). Note that, once the qualities and appeals have been recovered (using, say, Bernouilli sampling as suggested in [2]), it is easy to decide whether to use social influence when $r < 1$: Simply compare the expected number of purchases in both settings, using the equilibrium for the social influence case.

The Quality Ranking It is also important to emphasize that the quality ranking is not necessarily optimal asymptotically when $r < 1$, in contrast to the case where $r = 1$. However, the best possible equilibrium (in terms of the expected number of purchases) can be computed in strongly polynomial time using the reduction to the linear fractional assignment problem proposed in [2].

Recovering Product Quaity The results in this paper also provide an alternative way to estimate product qualities: Simply choose an arbitrary static rank and run the market with a sublinear signal until an equilibrium is reached. By Theorem 5.3, the market share of product i is

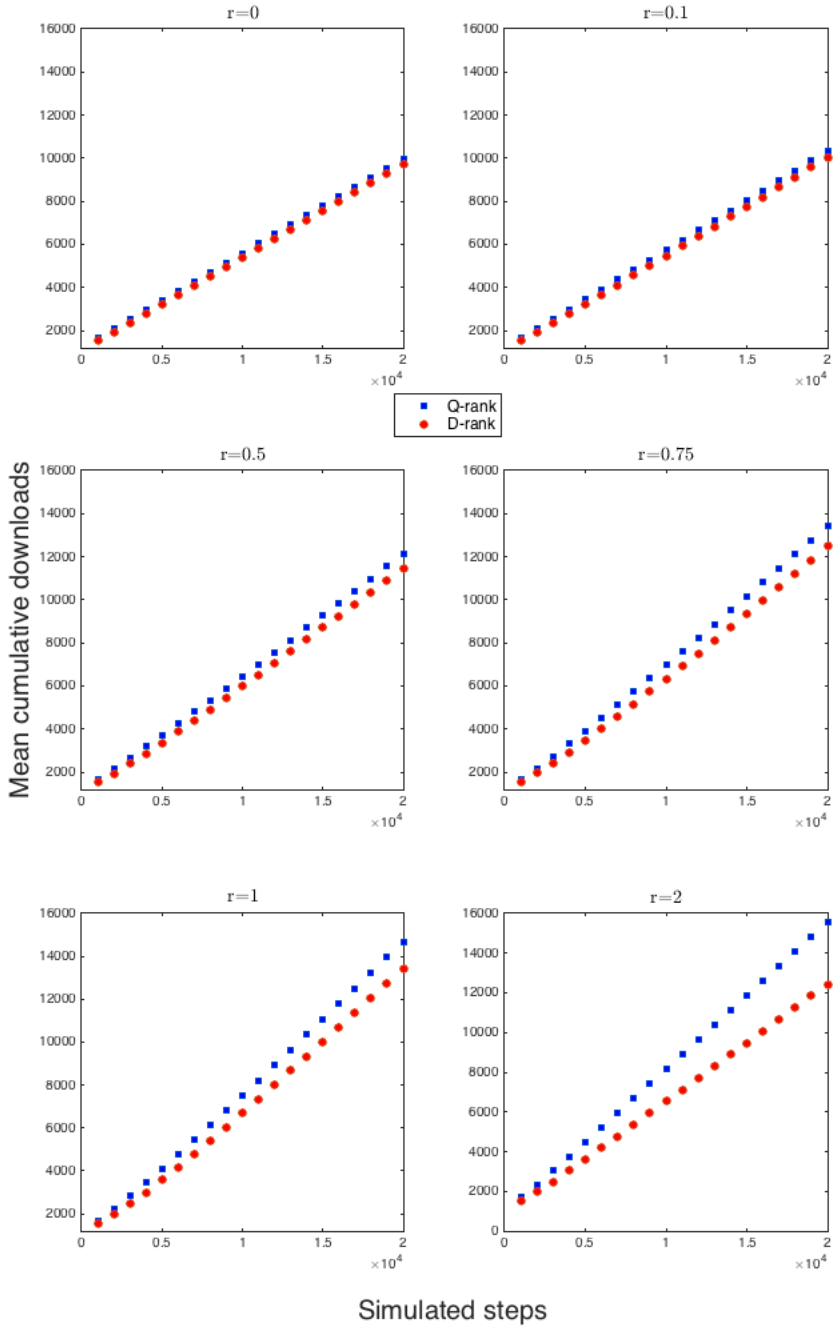

Figure 8: The Average Number of Downloads over Time for the Quality and Popularity Rankings for Various Social Signals in the First Setting for Song Appeal and Quality.

proportional to $v_i q_i$, which allows us to recover the quality of the products. This is only valid for sublinear signals obviously.

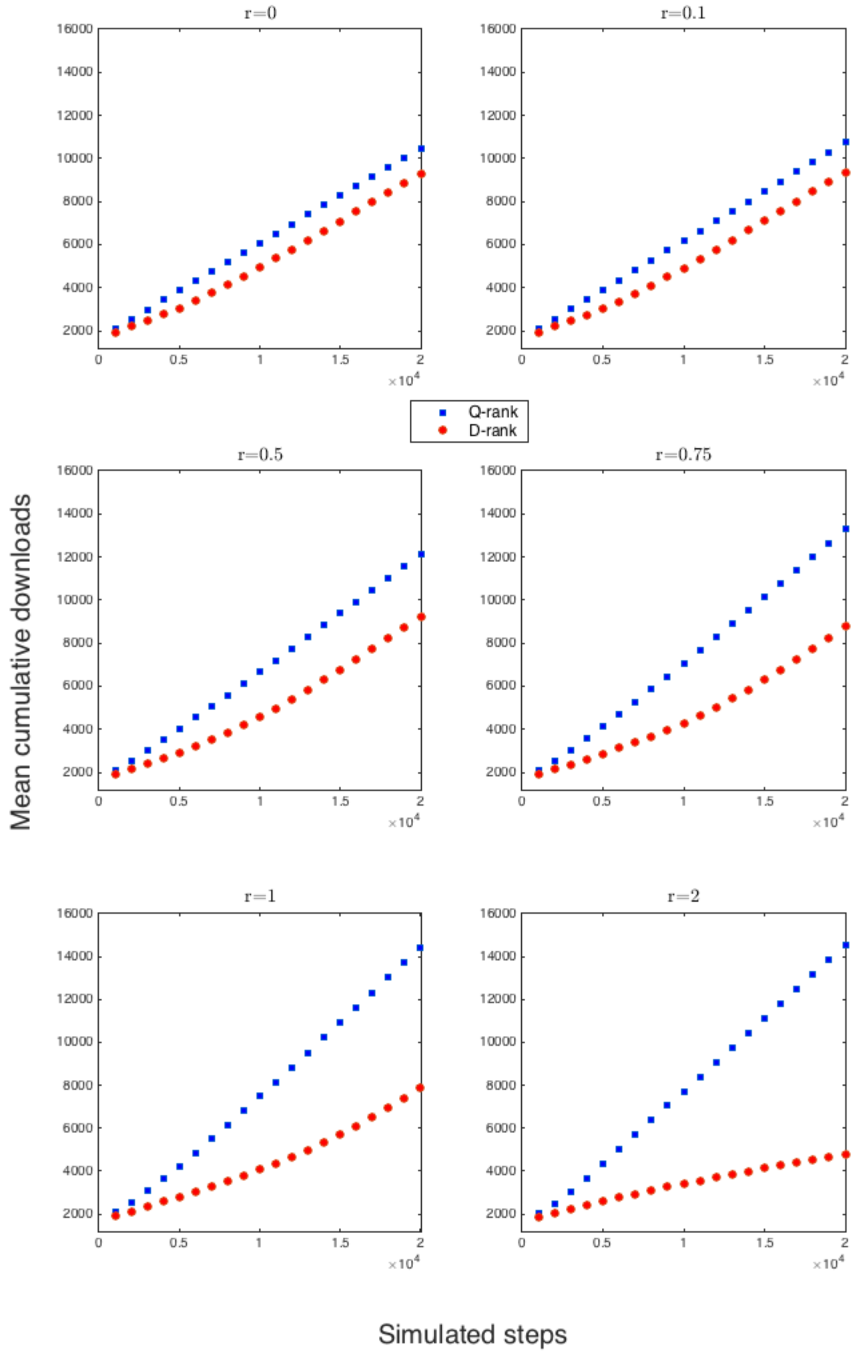

Figure 9: The Average Number of Downloads over Time for the Quality and Popularity Rankings for Various Social Signals in the Second Setting for Song Appeal and Quality.

8 Discussion and Conclusion

This paper studied the role of social influence in trial-offer markets where customer preferences are modeled by a generalization of a multinomial logit. In this model, both position bias and social influence impact the products tried by consumers.

The main result of the paper is to show that trial-offer markets, when the ranking of the products is fixed, converge to a unique equilibrium for sublinear social signals of the form d_i^r (e.g., display the square root of the number of downloads), where d_i represents the downloads of product i. Of particular interest is the fact that the equilibrium does not depend on the initial conditions, e.g., the product appeals, but only depends on the product qualities. Moreover, when the products are ranked by quality, i.e., the best products are assigned the highest visibilities, the equilibrium is such that the better products receive the largest market shares, which increase as r increases for the best products. The equilibrium for a sublinear social signal contrasts with the case with $r = 1$, where the market goes to a monopoly for the highest quality product (under the quality ranking). In the sublinear case, the market shares reflect product quality but no product becomes a monopoly. The paper also shows that, when $r > 1$, the market becomes more unpredictable. In particular, the inner equilibrium, which assigns a positive market share to all products, is unstable and the market almost surely converges to a monopoly. However, this monopoly depends on the initial conditions.

Simulation results on a setting close to the original MusicLab complemented the theoretical results. They show that the market converges quickly to the equilibrium for a sublinear social signal and that the convergence speed depends on the social signal strength. The simulation results also illustrate how the market shares of the highest (resp. lowest) quality products increase (resp. decrease) with r. As expected, when $r \leq 1$, the market is shown to be highly predictable, while it exhibits a lot of randomness when $r > 1$. The simulation results also show the benefits of social influence for market efficiency, i.e., the total number of purchases, and demonstrate that the quality ranking once again outperforms the popularity ranking.

Overall, these results shed a new light on the role of social influence in trial-offer markets and provide a comprehensive overview of the choices and tradeoffs available to firms interested in optimizing their markets with social influence. In particular, they show that social influence does not necessarily make markets unpredictable and is typically beneficial when the social signal is not too strong. Moreover, ranking the products by quality appears to be a much more effective policy than ranking products by popularity which may induce unpredictability and market inefficiency. The results also show that sublinear social signals give decision makers the ability to trade market efficiency for more balanced market shares.

Perhaps a main message of this paper is the sensitivity of the market to various design choices. The findings in [20] used the popularity ranking, which significantly affected their conclusions about market unpredictability and efficiency. The theoretical and simulation results of this paper, together with those in [2, 26] for the case $r = 1$, show that the market is highly predictable when using any static ranking and $r \leq 1$. Moreover, the quality ranking is optimal asymptotically when $r = 1$ and dominates the popularity ranking in all our simulations which were modeled after the MusicLab. This does not diminish the value of the results by Salganik et al [20] who isolated potential pathologies linked to social influence. But this paper shows that these pathologies are not inherent to the market but are a consequence of specific design choices in the experiment: The strength of the social signal and the ranking policy. Interestingly, it is only for a linear social signal that social influence can be shown to be always beneficial in expectation. Fortunately, for sublinear social signals, we can determine a priori if social influence is beneficial, given the analytic form of the equilibrium.

There are at least two potential research directions following this paper that worth investigating. First, it would be extremely valuable to construct large-scale cultural market experiment, varying the strengths of the social signal to complement our simulation results. Second, it would be interesting to extend our results to other settings including assortment problems (where the firm can select not only how to rank products but also which ones should be shown) [3] and to classical cascade models with a social signal [28].

References

[1] Abeliuk A., G. Berbeglia, F. Maldonado, and P. Van Hentenryck. Asymptotic optimality of myopic optimization in trial-offer markets with social influence. In *In the 25th International Joint Conference on Artificial Intelligence (IJCAI-16), New York, NY, July 9-15*, 2016.

[2] Andrés Abeliuk, Gerardo Berbeglia, Manuel Cebrian, and Pascal Van Hentenryck. The benefits of social influence in optimized cultural markets. *PLoS ONE*, 10(4), 2015.

[3] Andres Abeliuk, Gerardo Berbeglia, Manuel Cebrian, and Pascal Van Hentenryck. Assortment optimization under a multinomial logit model with position bias and social influence. *4OR*, 14(1):57–75, 2016.

[4] Michel Benaïm. Dynamics of stochastic approximation algorithms. In *Seminaire de probabilites XXXIII*, pages 1–68. Springer, 1999.

[5] Vivek S Borkar and Sean P Meyn. The ode method for convergence of stochastic approximation and reinforcement learning. *SIAM Journal on Control and Optimization*, 38(2):447–469, 2000.

[6] Simla Ceyhan, Mohammad Mousavi, and Amin Saberi. Social influence and evolution of market share. *Internet Mathematics*, 7(2), 2011.

[7] F. Chung, S. Handjani, and D. Jungreis. Generalizations of polya's urn problem. *Annals of Combinatorics*, 7:141–153, 2003.

[8] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. An experimental comparison of click position-bias models. In *Proceedings of the 2008 International Conference on Web Search and Data Mining*, pages 87–94. ACM, 2008.

[9] Marie Duflo and Stephen S Wilson. *Random iterative models*, volume 22. Springer Berlin, 1997.

[10] Per Engstrom and Eskil Forsell. Demand effects of consumers' stated and revealed preferences. *Available at SSRN 2253859*, 2014.

[11] Ming Hu, Joseph Milner, and Jiahua Wu. Liking and following and the newsvendor: Operations and marketing policies under social influence. *Management Science*, 2015.

[12] Coco Krumme, Manuel Cebrian, Galen Pickard, and Sandy Pentland. Quantifying social influence in an online cultural market. *PloS one*, 7(5):e33785, 2012.

[13] Harold J Kushner and George Yin. *Stochastic approximation and recursive algorithms and applications*, volume 35. Springer Science & Business Media, 2003.

[14] Kristina Lerman and Tad Hogg. Leveraging position bias to improve peer recommendation. *PloS one*, 9(6):e98914, 2014.

[15] Andrew Lim, Brian Rodrigues, and Xingwen Zhang. Metaheuristics with local search techniques for retail shelf-space optimization. *Management Science*, 50(1):117–131, 2004.

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

[17] Hosam Mahmound. *Polya Urn Models*. Chapman & Hall/CRC Texts in Statistical Science, 2008.

[18] Lev Muchnik, Sinan Aral, and Sean J. Taylor. Social influence bias: A randomized experiment. *Science*, 341(6146):647–651, 2013.

[19] Henrik Renlund. Generalized Polya Urns Via Stochastic Approximation. *ArXiv e-prints 1002.3716*, February 2010.

[20] 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.

[21] 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.

[22] 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.

[23] Christian Stummer, Elmar Kiesling, Markus Günther, and Rudolf Vetschera. Innovation diffusion of repeat purchase products in a competitive market: an agent-based simulation approach. *European Journal of Operational Research*, 245(1):157–167, 2015.

[24] Catherine Tucker and Juanjuan Zhang. How does popularity information affect choices? a field experiment. *Management Science*, 57(5):828–842, 2011.

[25] 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.

[26] P. Van Hentenryck, Abeliuk A., F. Berbiglia, G. Berbeglia, and F. Maldonado. Aligning popularity and quality in online cultural markets. In *In the Proceedings of the International AAAI Conference on Web and Social Media (ICWSM 2016), Cologne, Germany, May*, 2016.

[27] Giampaolo Viglia, Roberto Furlan, and Antonio Ladrón-de Guevara. Please, talk about it! when hotel popularity boosts preferences. *International Journal of Hospitality Management*, 42:155–164, 2014.

[28] Chi Wang, Wei Chen, and Yajun Wang. Scalable influence maximization for independent cascade model in large-scale social networks. *Data Mining and Knowledge Discovery*, 25(3):545–576, 2012.

[29] Xuchuan Yuan and H Brian Hwarng. Stability and chaos in demand-based pricing under social interactions. *European Journal of Operational Research*, 253(2):472–488, 2016.