

Taming the Matthew Effect in Online Markets with Social Influence

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

Pascal Van Hentenryck
 University of Michigan
 pvanhent@umich.edu

Abstract

Social influence has been shown to create a Matthew effect in online markets, increasing inequalities and leading to “winner-take-all” phenomena. Matthew effects have been observed for numerous market policies, including when the products are presented to consumers by popularity or quality. This paper studies how to reduce Matthew effects, while keeping markets efficient and predictable when social influence is used. It presents a market strategy based on randomization and segmentation, that ensures that the best products, if they are close in quality, will have reasonably close market shares. The benefits of this market strategy is justified both theoretically and empirically and the loss in market efficiency is shown to be acceptable.

Introduction

Salganik, Dodds, and Watts (2006) wrote a seminal paper on the negative influence of social influence in cultural markets. They created an experimental virtual market, called the MUSICLAB, in which participants can listen to songs and then download them if they like them. The songs are organized in a list or matrix form, giving different visibilities to the various songs, as is typically the case in online advertisement, online stores, or physical retail stores (e.g., (Craswell et al. 2008; Lim, Rodrigues, and Zhang 2004)). Each song was also associated with a popularity signal (e.g., (Engstrom and Forsell 2014; Viglia, Furlan, and Ladrón-de Guevara 2014)), i.e., the number of downloads of the song by earlier market participants. In addition, all the songs were ranked by popularity, which means that the song with the most downloads received the most visible position (i.e., the top of the list), the second most popular song the second most visible position, and so on. Salganik, Dodds, and Watts (2006) showed that social influence creates unpredictable markets, as well as significant Matthew effects (Rigney 2010). Indeed, multiple realizations of the MUSICLAB with various sets of participants show significantly different outcomes and strong inequalities in market shares between the songs. Those results have been reproduced by many authors in subsequent experiments or simulations (e.g., (Hu, Milner, and Wu 2016; Lerman and Hogg 2014; Muchnik, Aral, and Taylor 2013; van de Rijt et al. 2014)). The MUSICLAB is an example of

trial and offer markets which are particularly relevant in online settings. The popularity ranking is also widely used by firms running such markets in practice.

In recent years, some authors have started revisiting some of the assumptions in the MUSICLAB experiment and, in particular, the ranking policy which is the main controllable action of the firm running such a market. Van Hentenryck et al. (2016) showed that, when the products are ranked by quality instead of popularity, the market becomes asymptotically predictable and optimal: It leads a monopoly for the product of highest quality. Abeliuk et al. (2016) also proved that the dynamic ranking that optimizes the number of downloads at each step, which they called the performance ranking (Abeliuk et al. 2015), leads to the same outcome. These two ranking policies address the unpredictability and inefficiency of the market under social influence identified by Salganik, Dodds, and Watts (2006). However, they fail to address the market inequalities created by social influence: Two products of similar appeals and qualities will have fundamentally different outcomes with one becoming a monopoly in the long run. This “winner-takes-all” phenomena, although optimal from an efficiency standpoint, is typically considered undesirable.

This paper proposes a novel strategy that aims at addressing the three problems identified by Salganik, Dodds, and Watts (2006) simultaneously: unpredictability, inefficiencies, and inequalities. The strategy is a randomized segmentation protocol and is simple to deploy in online settings. The paper also analyzes its properties both theoretically and experimentally and shows that the protocol dramatically reduces inequalities among the best products in the market, while preserving high predictability and efficiency. The rest of the paper briefly reviews trial and offer markets, presents the new randomized segmentation protocol, and analyzes its properties both theoretically and experimentally.

Trial and Offer Markets

Following Krumme et al. (2012), this paper considers trial and offer markets modeled with an extended multinomial logit that includes product visibilities and social influence. In these markets, an incoming consumer observes n products in a list and a social influence vector d^t , where d_i^t is the number of purchases of product i at time t . She can then sample one product (e.g., listening to a song) and then de-

cides whether she wants to purchase it. Each position j in the list is associated with a visibility v_j that represents the inherent probability of sampling a product in position j or, intuitively, how much customers are attracted by a given position in the list. Each product i is characterized by two values: an appeal a_i which represents the probability of sampling product i and a quality q_i which captures the probability of purchasing product i after it has been sampled. Note that the appeals can capture some form of externalities such as marketing campaigns.

The firm running the market controls how to present the products to consumers, i.e., where the products are positioned in the list. In particular, the firm must choose a ranking θ that assigns a position $\theta(i)$ to each product i . As a result, the probability that a customer samples product i given ranking θ and social signal d is given by

$$p_i(\theta, d) = \frac{v_{\theta(i)}(a_i + d_i)}{\sum_{j=1}^n v_{\theta(j)}(a_j + d_j)}.$$

The probability that a customer purchases product i is then given by $p_i(\theta, d) q_i$. The firm is interested in maximizing the total number of purchases over time, i.e.,

$$\sum_t \sum_{i=1}^n p_i(\theta^t, d^t) q_i$$

where θ^t and d^t are the ranking and social signals at time t .

Prior work by Van Hentenryck et al. (2016) has shown that the quality ranking, a static ranking which ranks the products of highest quality to the highest visibilities (i.e., $q_i > q_j \Rightarrow v_{\theta(i)} \geq v_{\theta(j)}$), leads to markets where the product of highest quality becomes a monopoly asymptotically.¹ The same outcome holds for the so-called performance ranking that determines the optimal ranking θ^* at each step t , i.e., $\theta^* = \arg\text{-max}_{\theta} \sum_{i=1}^n p_i(\theta, d^t) q_i$. These two ranking policies are optimal and predictable asymptotically but they obviously create major Matthew effects: *A small difference in quality leads to a product losing its entire market share in the long run.*

Randomized Segmentation

We now propose a Randomized Segmentation Protocol (RSP) to address the Matthew effect, while preserving predictability and efficiency. The RSP is presented in Figure 1: It uses two key ideas to tame the Matthew effect. First, it *segments* the market in m submarkets that we call *worlds*. These worlds evolve independently with their own social influence signals and ranking policies. This segmentation itself has no impact on the Matthew effect if the quality or performance ranking is used, since the best-quality product will still obtain a monopoly in each world asymptotically and a large market share over a finite horizon. To counteract the Matthew effect, the RSP perturbs the product qualities slightly with some random noise of zero means and then applies the quality ranking on these perturbed qualities.

¹If several products have the same highest product $v_{\theta(i)} q_i$, then the market becomes a beta distribution among them.

Overview of The Randomized Segmentation Protocol

1. The market is organized in m sub-markets called worlds.
 2. In each world, the product qualities are perturbed with a random noise of zero mean. The products are displayed using the quality ranking over the perturbed qualities.
 3. Each world maintains social influence signals independently of the other world.
 4. When a customer enters the market, she is randomly assigned to a world and presented with the ranking and the social signals of that world.
 5. If the customer buys the product she sampled, the social signal of that product in that world is increased.
-

Figure 1: The Randomized Segmentation Protocol (RSP).

The qualities obviously remain the same for the purpose of downloading the products. As a result, each world may now use a different ranking and a different product may now take a substantial market share. Indeed, if the product with the second highest quality is placed in first position, it may become the most attractive product through the position visibilities, i.e., $v_1 q_2 > v_2 q_1$. As a consequence, the RSP will balance the market shares across the different worlds, ensuring a fairer distribution of the market shares.

Theoretical Analysis

This section formalizes the RSP more precisely and presents a theoretical analysis of its behavior both in terms of inequalities, efficiency, and predictability. Given the complexity of the formulas, the analysis is only carried out for the perturbations of the first three products, but it generalizes naturally to more products. For simplicity, assume that the product qualities satisfy $q_1 \geq \dots \geq q_n$ and that the qualities of the top three products are reasonably close, i.e.,

$$q_2 = q_1 - \epsilon_2 \tag{1}$$

$$q_3 = q_1 - \epsilon_3 \tag{2}$$

where $\epsilon_3 \geq \epsilon_2 \geq 0$. More precisely, we assume that ϵ_2 and ϵ_3 are small enough such that $v_1 q_3 > v_2 q_1$. As a result, a ranking $(3, 1, 2)$ ensures that product 3 has the best normalized quality (i.e., its inherent quality times the visibility of its position) and the market will go to a monopoly for product 3 asymptotically.

Our proposed randomized segmentation protocol partitions the market in m worlds and defines perturbed qualities $\hat{q}_{i,w}$ for each product i and world w as

$$\hat{q}_{i,w} = q_i + \eta_{i,w}$$

where $\eta_{i,w}$ is a normally distributed noise with zero mean and standard deviation σ . In other words, $\eta_{i,w}$ is a random variable distributed under the normal distribution

$$f_{\sigma}(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right) \tag{3}$$

whose CDF $F_\sigma(x)$ is given by

$$F_\sigma(x) = \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{x}{\sigma\sqrt{2}} \right) \right].$$

Combining this with the fact that the highest qualities are close, the perturbed qualities can thus be rewritten as

$$\begin{aligned}\hat{q}_{1,w} &= q_1 + \eta_{1,w}, \\ \hat{q}_{2,w} &= q_1 - \epsilon_2 + \eta_{2,w}, \\ \hat{q}_{3,w} &= q_1 - \epsilon_3 + \eta_{3,w}\end{aligned}$$

In each world w , the products are ranked in terms of their perturbed qualities $\hat{q}_{i,w}$ instead of their intrinsic quality q_i . Hence, depending on the outcomes of the random variables $\eta_{1,w}$, $\eta_{2,w}$ and $\eta_{3,w}$, the products may be ranked in the sequences (1,2,3), (1,3,2), (2,1,3), (2,3,1), (3,1,2), or (3,2,1). As a result, asymptotically, the different worlds will converge to different monopolies: In some of them, the market will become a monopoly for product 1 while, in other worlds, the market will become a monopoly for product 2 or product 3. To derive the theoretical results, we will use the following notations:

$$\begin{aligned}x_i^+ &\doteq x + \epsilon_i, \\ x_i^- &\doteq x - \epsilon_i, \\ \eta_{i,w}^+ &\doteq \eta_{i,w} + \epsilon_i \text{ and} \\ \eta_{i,w}^- &\doteq \eta_{i,w} - \epsilon_i \text{ for } i \in \{2, 3\}.\end{aligned}$$

Theorem 1. *When the perturbed quality ranking is performed in m different worlds, the number of worlds in which products 1, 2 and 3 are ranked first (denoted by $s_{1,m}$, $s_{2,m}$ and $s_{3,m}$ respectively) are given by the following binomial distributions*

$$\mathbb{P}(s_{1,m} = k) = \binom{m}{k} P_1^k (P_2 + P_3)^{m-k}$$

$$\mathbb{P}(s_{2,m} = k) = \binom{m}{k} P_2^k (P_1 + P_3)^{m-k}$$

$$\mathbb{P}(s_{3,m} = k) = \binom{m}{k} P_3^k (P_1 + P_2)^{m-k}$$

where

$$P_1 = \int_{-\infty}^{\infty} f_\sigma(x) F_\sigma(x_2^+) F_\sigma(x_3^+) dx \quad (4)$$

$$P_2 = \int_{-\infty}^{\infty} f_\sigma(x_2^+) F_\sigma(x) F_\sigma(x_3^+) dx \quad (5)$$

$$P_3 = \int_{-\infty}^{\infty} f_\sigma(x_3^+) F_\sigma(x) F_\sigma(x_2^+) dx. \quad (6)$$

Proof. Let P_1 be the probability that $\hat{q}_{1,w}$ be greater than $\max(\hat{q}_{2,w}, \hat{q}_{3,w})$, i.e.,

$$\begin{aligned}P_1 &= \mathbb{P}(\hat{q}_{1,w} > \max(\hat{q}_{2,w}, \hat{q}_{3,w})) \\ &= \mathbb{P}(\eta_{1,w} > \max(\eta_{2,w}^-, \eta_{3,w}^-)) \\ &= \int_{-\infty}^{\infty} f_\sigma(x) \mathbb{P}(x > \eta_{2,w}^-) \mathbb{P}(x > \eta_{3,w}^-) dx \\ &= \int_{-\infty}^{\infty} f_\sigma(x) F_\sigma(x_2^+) F_\sigma(x_3^+) dx.\end{aligned}$$

Now let P_2 be the probability that $\hat{q}_{2,w}$ be greater than $\max(\hat{q}_{1,w}, \hat{q}_{3,w})$, i.e.,

$$\begin{aligned}P_2 &= \mathbb{P}(\hat{q}_{2,w} > \max(\hat{q}_{1,w}, \hat{q}_{3,w})) \\ &= \mathbb{P}(\eta_{2,w}^- > \max(\eta_{1,w}, \eta_{3,w}^-)) \\ &= \int_{-\infty}^{\infty} f_\sigma(x_2^+) \mathbb{P}(x > \eta_{1,w}) \mathbb{P}(x > \eta_{3,w}^-) dx \\ &= \int_{-\infty}^{\infty} f_\sigma(x_2^+) F_\sigma(x) F_\sigma(x_3^+) dx.\end{aligned}$$

Finally, let P_3 be the probability that $\hat{q}_{3,w}$ be greater than $\max(\hat{q}_{1,w}, \hat{q}_{2,w})$.

$$\begin{aligned}P_3 &= \mathbb{P}(\hat{q}_{3,w} > \max(\hat{q}_{1,w}, \hat{q}_{2,w})) \\ &= \mathbb{P}(\eta_{2,w}^- > \max(\eta_{1,w}, \eta_{3,w}^-)) \\ &= \int_{-\infty}^{\infty} f_\sigma(x_3^+) \mathbb{P}(x > \eta_{1,w}) \mathbb{P}(x > \eta_{2,w}^-) dx \\ &= \int_{-\infty}^{\infty} f_\sigma(x_3^+) F_\sigma(x) F_\sigma(x_2^+) dx\end{aligned}$$

Now the number of times $s_{i,m}$ that product i is ranked first among m worlds is binomially distributed since, in each world, there is a probability P_i that product i is ranked first.

$$\mathbb{P}(s_{1,m} = k) = \binom{m}{k} P_1^k (P_2 + P_3)^{m-k}$$

$$\mathbb{P}(s_{2,m} = k) = \binom{m}{k} P_2^k (P_1 + P_3)^{m-k}$$

$$\mathbb{P}(s_{3,m} = k) = \binom{m}{k} P_3^k (P_1 + P_2)^{m-k}.$$

□

Corollary 2. *The expectation and variance of $s_{i,m}$ for $i \in \{1, 2, 3\}$, is given by*

$$\mathbb{E}[s_{i,m}] = mP_i$$

$$\operatorname{Var}[s_{i,m}] = mP_i(1 - P_i).$$

Furthermore, the probabilities that the $s_{i,m}$'s fall between their expected value $\pm X\%$ are given by:

$$\begin{aligned}\mathbb{P} \left(|s_{1,m} - \mathbb{E}[s_{1,m}]| \leq \frac{X}{100} \mathbb{E}[s_{1,m}] \right) \\ = \sum_{k=\lceil (1-\frac{X}{100})\mathbb{E}[s_{1,m}] \rceil}^{\lfloor (1+\frac{X}{100})\mathbb{E}[s_{1,m}] \rfloor} \binom{m}{k} P_1^k (P_2 + P_3)^{m-k},\end{aligned} \quad (7)$$

$$\begin{aligned}\mathbb{P} \left(|s_{2,m} - \mathbb{E}[s_{2,m}]| \leq \frac{X}{100} \mathbb{E}[s_{2,m}] \right) \\ = \sum_{k=\lceil (1-\frac{X}{100})\mathbb{E}[s_{2,m}] \rceil}^{\lfloor (1+\frac{X}{100})\mathbb{E}[s_{2,m}] \rfloor} \binom{m}{k} P_2^k (P_1 + P_3)^{m-k}\end{aligned} \quad (8)$$

$$\begin{aligned}\mathbb{P} \left(|s_{3,m} - \mathbb{E}[s_{3,m}]| \leq \frac{X}{100} \mathbb{E}[s_{3,m}] \right) \\ = \sum_{k=\lceil (1-\frac{X}{100})\mathbb{E}[s_{3,m}] \rceil}^{\lfloor (1+\frac{X}{100})\mathbb{E}[s_{3,m}] \rfloor} \binom{m}{k} P_3^k (P_1 + P_2)^{m-k}.\end{aligned} \quad (9)$$

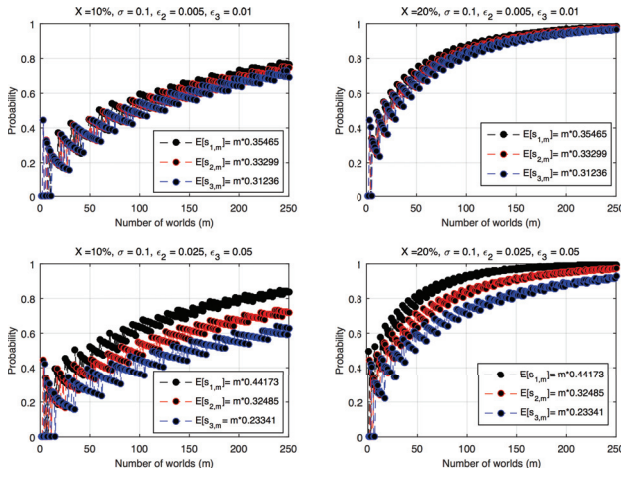


Figure 2: Probabilities that the number of worlds in which products 1, 2 and 3 become a monopoly fall in the expected value $\pm X\%$ of the total number of worlds, as a function of the number of worlds.

We now quantify the loss ratio in purchases induced by the RSP compared to the (asymptotically) optimal quality ranking, i.e.,

$$l_t \doteq \frac{d_q^t - d_{rsp}^t}{d_q^t}, \quad (10)$$

where d_q^t is the total number of purchases under the quality ranking at time t and d_{rsp}^t is the total number of purchases of the RSP at t . Expression (10) can be approximated asymptotically with

$$d_q^t \sim q_1 t \quad (11)$$

and

$$d_{rsp}^t \sim (q_1 P_1 + q_2 P_2 + q_3 P_3) t, \quad (12)$$

Equations (11) and (12) together with definition (10) can be used to characterize the asymptotic loss ratio, i.e.,

$$l_t^\infty = \frac{\epsilon_2 P_2 + \epsilon_3 P_3}{q_1}. \quad (13)$$

Observe that, when q_2 and q_3 are equal to q_1 , the RSP is optimal but has no Matthew effect contrary to the quality and performance ranking.

We now investigate the consequences of these results. Figure 2 plots the probability that the number of worlds in which product i becomes a monopoly falls in between its expected value $\pm X\%$ as a function of the number of worlds. The plots are obtained from Equations 7–9 for various values of parameters ϵ_2 and ϵ_3 given $\sigma = 0.1$. The figure also reports the formula for the expected market shares that show a nice distribution of the purchases among the products. The plots also show a rapid convergence around the mean as the number of worlds increases.

Figure 3 depicts the probabilities that a given product becomes a monopoly as a function of ϵ_2 assuming that $\epsilon_3 = 2\epsilon_2$ and $\sigma = 0.1$. The results show that, as the distance in quality increases, the market shares also grow apart




Figure 3: Probabilities that products 1, 2, and 3 become a monopoly as a function of ϵ_2 where $\epsilon_3 = 2\epsilon_2$ and $\sigma = 0.1$.




Figure 4: Probability that product 1, 2 and 3 become a monopoly as a function of σ where $\epsilon_3 = 2\epsilon_2 = 0.01$.

but they remain nicely distributed among the products as desired.

Figure 4 depicts the probabilities that a given product becomes a monopoly as a function of the standard deviation $\sigma = 0.1$ given $\epsilon_2 = 0.005$ and $\epsilon_3 = 2\epsilon_2$. The figure shows how to use the standard deviation to balance the market shares between the products. It shows that small standard deviations already balance the market shares well, opening significant opportunities for controlling the market towards a desired outcome.

Figures 5 plots the loss ratio as a function of ϵ_2 . The ratio initially increases as ϵ_2 increases, since a lower quality product may become a monopoly in some worlds. Once ϵ_2 approaches the standard deviation, the loss ratio starts to decrease since this sub-optimal behavior is increasingly less likely to occur.

Experimental Results

We now describe experimental results on the MusicLab settings (Salganik, Dodds, and Watts 2006). Like previous stud-




Figure 5: The Loss Ratio (13) as a Function of ϵ_2 When $\epsilon_3 = 2\epsilon_2$ and $\sigma = 0.1$.




Figure 6: The Inequality Measure (14) over 50 Experiments. The largest possible Matthew effect is $\sqrt{2}$.

ies, the experiments use the setting in (Krumme et al. 2012) which gives specific values for appeals, qualities, and visibilities. The qualities of the second and third best products are modified to obtain a setting in which $\epsilon_2 = 0.005$ and $\epsilon_3 = 0.01$ (the quality of the best product is $q_1 = 0.8$). This makes the Matthew effect particularly dramatic, since a small difference in quality induces large market inequalities. The experiments aim at confirming the theoretical results within a finite horizon and at balancing the market shares between the top three products. The experiments also use a standard deviation of $\sigma = 0.1$. The presentation discusses the effectiveness of the RSP to tame the Matthew effect and the cost involved in doing so. We also present some results obtained on the original MUSICLAB data for completeness.

Taming the Matthew Effect Figure 6 depicts the effectiveness of the RSP to tame the Matthew effect. It reports




Figure 7: Number of Purchases of the First 3 Products.




Figure 8: Number of Purchases For the First 3 Products.

the inequality measure

$$\frac{\sqrt{(d_1^t - d_2^t)^2 + (d_1^t - d_3^t)^2 + (d_3^t - d_2^t)^2}}{d_1^t + d_2^t + d_3^t} \quad (14)$$

for the quality ranking and the RSP as the number of customers increases. The plot reports the mean (solid) and the mean \pm the variance (dashed lines) for 50 simulations. The largest possible Matthew effect with this measure is $\sqrt{2}$. The figure shows that the quality ranking converges towards 1 in average, while the RSP converges towards a value around 0.45. This represents a dramatic reduction in inequalities among the best products.

Figure 7 presents box plots for the total purchases of the top three products and clearly explains why there is such a reduction in the Matthew effect. The figures summarize 50 experiments, each with 2 million customers. They show that the quality ranking may produce significant Matthew effects since the third quartile exhibits purchases up to the order of $13 \cdot 10^5$ and $11 \cdot 10^5$ for the first two products. This severe Matthew effect arises despite the fact that the products have almost the same quality, showing the strong need of taming this effect in practice to obtain reasonably fair markets. In contrast, the RSP has a third quartile in the order of $8 \cdot 10^5$ and $6 \cdot 10^5$ for the same products. Observe also the higher medians and the smaller variances in the RSP.

Figure 8 depicts the total number of purchases for the first three products in the quality ranking and in the RSP. The plots show that the quality ranking has more aggregate purchases. We will come back to this when we discuss the cost of fairness, i.e., the cost of taming the Matthew effect. Note also that the variance also increases in the RSP when aggregating the purchases of the three products, which is due to the added Gaussian noise and the weaker social influence since the market is split in several words. Of course, the variance for each product independently decreases in a substantial way, as discussed earlier.

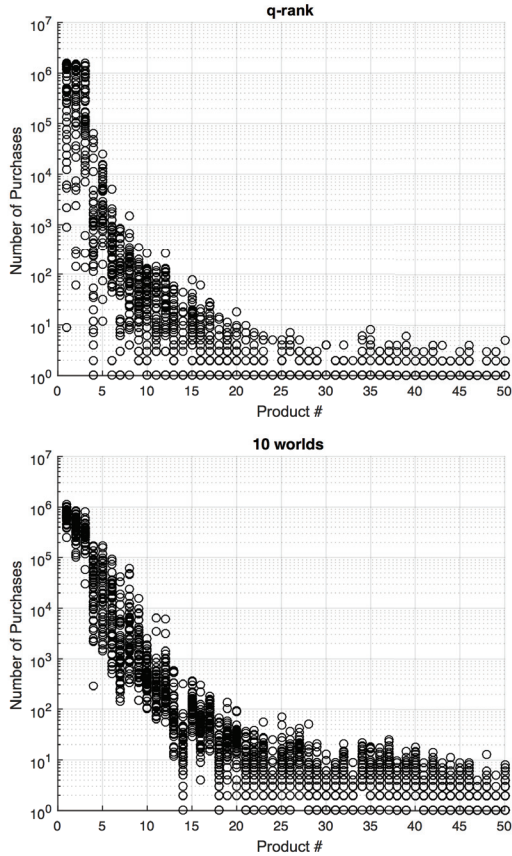


Figure 9: Purchase distribution for 500 experiments with 2 million customers for the quality ranking and the RSP with 10 worlds (right).

Figure 9 describes an interesting side-effect of the RSP: a significant decrease in outliers and variances for the total purchases of all products. Observe how the top three products (left side of the bottom plot) now have very similar distributions and how the products in the range 4–15 have much smaller variances. The segmentation of the market has strong benefits on the predictability of the market as well.

The Cost of Fairness Intuitively, the RSP has two net effects that make the market less efficient. First, it reduces the impact of social influence which is spread over many worlds, increasing the randomness in the early stage of the market. Second, it allows the market to achieve monopolies that are not optimal (i.e., monopolies of products whose quality is not the best among all products). Figure 10 quantifies this loss of efficiency: It depicts the average loss in purchases when using the RSP with 10 worlds. As can be seen, the loss is small and decreases significantly as the number of customers increases, matching the theoretical predictions. Figure 12 depicts the loss for the actual MUSICLAB data for various values of σ and the noise added to all songs. Figure 12 shows how the loss varies at the number of worlds increases for different number of customers. The figure shows




Figure 10: Purchase Loss Ratio (10 Worlds).

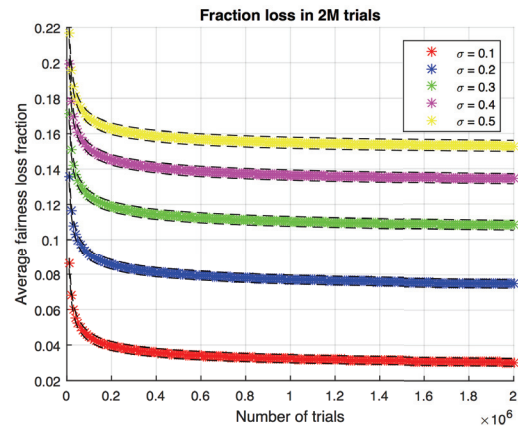


Figure 11: Purchase Loss Ratio (10 Worlds) on the Actual MusicLab Data.

that the loss slowly increases with the number of worlds.

Conclusion

This paper presented a randomized segmentation protocol (RMS) to jointly address the unpredictability, inefficiency, and inequalities created by social influence in trial and offer markets. The RMS was shown, both through theoretical analysis and simulations, to remove the Matthew effect among the best products, balancing the market shares as a function of product qualities. The RMS also leads to predictable and efficient markets as the loss in efficiency is shown to be very small and tends to zero as the product qualities tend to the same values. As a result, the RMS seems to be an appealing, easy to deploy marketing strategy to leverage social influence in trial and offer markets.

References

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

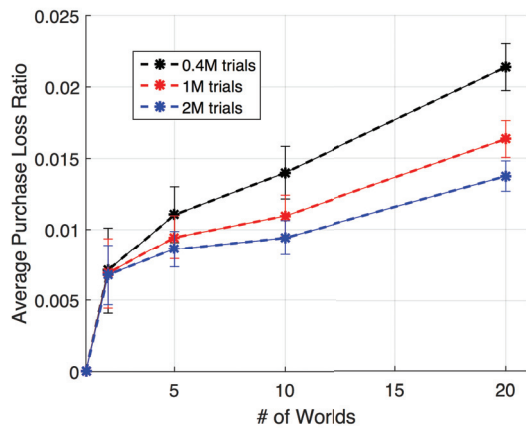


Figure 12: Purchase loss ratio as a function of the number of worlds for several number of trials. The plot shows the average of 100 experiments and the variance is shown with the error bars. The parameters chosen where $\sigma = 0.1$, $\epsilon_3 = 2\epsilon_2 = 0.01$.

Abeliuk, A.; Berbeglia, G.; Maldonado, F.; and Hentenryck, P. V. 2016. Asymptotic optimality of myopic optimization in trial-offer markets with social influence. In Kambhampati, S., ed., *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, 2458–2464. IJCAI/AAAI Press.

Craswell, N.; Zoeter, O.; Taylor, M.; and Ramsey, B. 2008. An experimental comparison of click position-bias models. In *Proceedings of the 2008 International Conference on Web Search and Data Mining*, 87–94. ACM.

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

Hu, M.; Milner, J.; and Wu, J. 2016. Liking and following and the newsvendor: Operations and marketing policies under social influence. *Management Science* 62(3):867–879.

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):e98914.

Lim, A.; Rodrigues, B.; and Zhang, X. 2004. Metaheuristics with local search techniques for retail shelf-space optimization. *Management Science* 50(1):117–131.

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

Rigney, D. 2010. *The Matthew Effect*. Columbia University Press.

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.

van de Rijt, A.; Kang, S. M.; Restivo, M.; and Patil, A. 2014. Field experiments of success-breeds-success dynam-

ics. *Proceedings of the National Academy of Sciences* 111(19):6934–6939.

Van Hentenryck, P.; Abeliuk, A.; Berbeglia, F.; Maldonado, F.; and Berbeglia, G. 2016. Aligning popularity and quality in online cultural markets. *Tenth International AAAI Conference on Web and Social Media (ICWSM-16)*.

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.