

**Blockbuster culture's next rise or fall:
The effect of recommender systems on sales diversity**

Daniel Fleder Kartik Hosanagar

The Wharton School, University of Pennsylvania, Philadelphia, PA

1 Introduction

Recommender systems offer benefits to both consumers and firms. For consumers, recommender systems help individuals both become aware of new products as well as select desirable products among myriad choices (Pham & Healey, 2005). For firms, recommender systems have the potential to increase profits by converting browsers into buyers, cross-selling products, and increasing loyalty through a customized browsing experience (Schafer et al., 1999). Major online firms, such as Amazon.com, Netflix, and Yahoo! Music, host recommender systems, and the prevalence of such systems is expected to grow.

In this paper, we study the effect of recommender systems on consumer purchasing of hedonic goods. Two different views exist on the impact of recommenders on diversity of sales. Some believe recommenders will help consumers discover new products (Brynjolffson et al. 2006). On the other hand, others believe the design of common recommender systems will only reinforce the popularity of already popular titles (Mooney & Roy, 1999). The first belief leads to a hypothesis that recommender systems will increase sales diversity. The second belief leads to a hypothesis that recommender systems will decrease sales diversity. This paper is a first attempt to reconcile these seemingly incompatible views.

We develop an analytical model of recommender systems and their influence on consumer choice. The model allows us to explore the strong, path-dependent interaction between recommendations and sales. We believe this paper is the first to attempt to isolate the impact of recommender systems on long-run sales diversity.

2 Prior work

Despite the growing use of recommender systems, little is known about how these systems affect consumer behavior. In marketing, Cooke et al. (2002) examine how the likelihood of purchasing a recommended product depends on the information provided, context, and familiarity with other products simultaneously recommended. Senecal et al. (2004) experimentally show that recommendations do influence choice. They find that online recommender systems can have even more influence than human recommendations, and the effect is often greater for experience goods than search goods.

To the best of our best knowledge, there have been no studies that have attempted to isolate the specific impact of recommenders on sales diversity. However, the topic has received mention from several researchers. For example, Mooney and Roy (1999) suggest that collaborative filters may perpetuate homogeneity in a library setting, but this brief comment is used more to suggest difficulties of working with sparse data than to examine social implications. Brynjolffson et al (2006) discuss the supply-side and demand-side drivers of increased sales diversity on the Internet. On the supply-side are factors such as the ability to stock more products (i.e., product diversity). On the demand-side are active search tools that make it easy for consumers to locate products on websites and passive tools such as recommender systems. Holding supply-side factors constant, the authors find that demand-side factors on the Internet contributed to greater sales diversity on the Internet versus catalog sales for a medium-sized retailer. However, the authors did not isolate the effects of recommenders. Prior work reveals two themes. One, much of the extant body of work has not tried to isolate the impact of recommenders on the diversity of sales. Two, there exist incompatible views on the impact of recommenders on long-run sales diversity.

3 Problem definition

We consider a market with a firm selling a single class of hedonic good. (e.g., music CDs). The dependent variable studied is sales diversity, which we define and measure by the Gini coefficient (Dorfman, 1979). The Gini coefficient is a common measure of distributional inequality in many fields. Let $L(u)$ be the Lorenz curve denoting the percentage of the firm's revenue generated by the lowest $100u\%$ of goods sold during a fixed time period, $u \in [0,1]$. The Gini coefficient is then defined as

$$G = 1 - 2 \int_0^1 L(u) du$$

Graphically, this is equivalent to $G = A/(A+B)$ (Fig 1). A value of $G = 0$ implies all products have equal sales. In contrast, values closer to 1 represent unequal distributions, in which a small fraction of products accounts for the majority of sales. This paper’s central question asks whether recommender systems will push the Lorenz curve one way or the other. As discussed above, anecdotal views do not yield consensus.




Figure 1. Graphical interpretation of the Lorenz curve and Gini coefficient

Denote by G_0 the Gini coefficient during a fixed time period in which a recommender system was *not* used. In contrast, let G_t be the Gini coefficient in which a recommender system was employed with all else equal. The recommender system is said to have a homogeneity bias, diversity bias, or no bias depending on the following conditions

$$\text{Bias} = \begin{cases} \text{Homogeneity bias,} & G_t > G_0 \\ \text{Diversity bias,} & G_t < G_0 \\ \text{Neutral,} & G_t = G_0 \end{cases}$$

Below, we test the bias induced by common recommenders that recommend products based on historical sales (e.g., Amazon’s collaborative filter). For modeling purposes, we use a stylized recommender.

4 Proposed Analytic model

We begin with a case in which a firm offers two horizontally differentiated products termed 1 and 0. Our firm has M identical consumers who make purchases in sequence. These customers have been identified to be of the same type based on the purchase of one of more outside goods. Without recommendations, consumers purchase product 1 and 0 with respective probabilities p and $\bar{p} = 1 - p$.¹ A recommender system’s input is the firm’s sales database. Define a sales history $H_t = \{I_1, \dots, I_t\}$ to be the record of which items $I_i \in \{0, 1\}$ were purchased on each occasion. An initial representation for the recommender system and its influence on behavior is the following two-part process.

$$S_t = f(H_{t-1}, p) \tag{1}$$

$$P_t = g(S_t, p) \tag{2}$$

The recommender system is a function f that maps a sales history H_{t-1} and customer profile p into a signal, or recommended product, $S_t \in \{0, 1\}$ shown to the consumer. Since we model a collaborative filter, the product’s intrinsic content can be ignored (as opposed to content-based systems, collaborative filters use sales data only). Next, an influence function g describes how the consumer responds to the recommendation: g maps the recommended product S_t and customer profile p into an updated purchase probability $P_t \equiv P(I_t = 1)$. Recognizing that f is a deterministic function, we can combine the equations into a single expression. Accordingly, a recommender system r is a function mapping the sales history and customer type into an updated purchase probability.

¹ All consumers are required to make some purchase, so the choice is between products and not whether to buy. Examining purchase incidence would also be interesting. We exclude it in order to isolate choice-specific effects.

$$P_t = r(H_{t-1}, p) \quad (3)$$

We next assume that all relevant information in H_t can be represented by two numbers t and $K_t = \sum_{i=1}^t I_i$. This statement means that at time t , the only important data is that there were K_t and $t - K_t$ purchases of products 1 and 0 respectively. This is reflective of collaborative filters in which time-order is not considered. With this property,

$$P_t = r(K_{t-1}, t - 1, p) \quad (4)$$

Before presenting the function r , we first give the intuition. Assume consumers are initially indifferent between the products so that $p = 1/2$. Suppose the first consumer purchases product 1. The recommender observes this and considers this purchase as an implicit vote for product 1. When the next customer arrives the recommender will be slightly more likely to recommend product 1; in turn, this means, since consumers' initial preferences are identical, that the second consumer will be slightly more likely to buy product 1 too. If consumer two buys product 1, this will end up increasing product 1's salience even more for the third customer. In contrast, if the second customer had bought product 0, the recommender would notice this and adjust accordingly. Mathematically, we can represent this scenario in a simple way. Our identical consumers flip coins to decide between products. Without recommendations, they would all flip the same coin p . With recommendations, however, the coin's bias can change over time to reflect the recommender's influence. Just how the bias changes, however, depends on past purchases. Let the function r be such that

$$P_t = \frac{a + K_{t-1}}{b + (t - 1)}, \text{ where } a \leq b \quad (5)$$

In particular, as our base case set $(a, b) = (1, 2)$ so that

$$P_t = \frac{1 + K_{t-1}}{1 + t} \quad (6)$$

At $t = 1$ a fair coin is used, and so $p_1 = 1/2$. Suppose the first purchase was product 1, so that $K_1 = 1$. Then $p_2 = (1 + 1)/(2 + 1) = 2/3$; that is, after observing a purchase of product 1, the recommender places more emphasis on this product, and hence the next consumer is more likely to buy it: $p_2 = 2/3 > p_1 = 1/2$. In contrast, suppose the first purchase had been product 0. Then we would have $K_1 = 0$, so that $p_2 = (1 + 0)/(2 + 1) = 1/3 < p_1$.

Theorem 1. The sequence $\{P_i\}$ converges to a random variable P with probability 1.

Theorem 2. At any t , P_t is uniformly distributed with $P(P_t = c) = \frac{1}{t}$ for $c = \frac{1}{t+1}, \dots, \frac{t}{t+1}$.

We skip the proofs here due to space constraints.

4.1 Interpreting model results

Theorems 1-2 indicate that the process always converges to a fixed market share. However, what that share will be, one cannot say in advance. The figure below graphically illustrates this effect (a similar illustration was observed by Arthur (1990)). The figure shows ten random runs from our model. Each run starts with $p = 0.5$ and lasts for 500 time periods. We see that each run does appear to converge to some limiting value. However, we also see that these 10 limiting values are very different. As is guaranteed by the theorem, they are equally likely to take on any limiting market share on the unit interval.

We can speculate that these effects are partly responsible for resurgence of the book *Touching the Void* as described by Anderson (2004). *Touching the Void* had been in print for some time and was experiencing declining sales. Amazon.com's recommendation algorithm somehow began associating it with a well-selling book, Krakauer's *Into Thin Air*. As soon as this occurred, *Touching the Void's* fortune reversed, and it became a best-seller. Based on the above results, we can see why a few chance purchases could lead to higher likelihood of it being recommended, which in turn may increase sales. As suggested by Figure 2, there can be strong path dependence in the long-run outcome.




Figure 2. Sample runs, each starting at $p = 0.5$.

Assesing net impact of the recommender.

Without recommendation, a $p = 0.5$ implies $E[G_0(P_{t=\infty})] = E[G_0(p)] = G_0(p) = 0$. In a world with the recommender, $E[G_1(P_{t=\infty})] = \frac{1}{4} > 0$ (proof omitted), showing the recommender has a homogeneity bias. The result has another interesting characteristic. Theorem 2 leads to a corollary: $E[P_t] = 0.5$ for any t . This means that averaged across all sample paths, neither product is favored a priori (i.e. the outcome can go in either direction, as seen by the symmetry in Figure 2). However, while we cannot say which product will win, we can say that some product is likely to win. That winner is likely to be the product with a greater number of chance successes early on. Thus while recommenders are often expected to increase diversity, the use of implicit votes, which is common in practice, could have the opposite effect of reinforcing already popular titles.

4.2 Model generalizations

This paper posited a “changing coin” model and derived its limiting distribution. In fact, a similar process called the Polya process has been studied rather well (Johnson & Kotz, 1977). The Polya process can be seen as a particular type of sampling with replacement. Consider an urn initially containing b black balls and w white balls. A ball is drawn at random. That ball is returned to the urn with s more balls of the same color. In terms of a recommender system, one can imagine the selection of a ball from the urn as consumer choice (i.e., realized sales). The s additional balls of the same color added to the urn model the fact that the recommender reacts to the realized sale by recommending the product more often and thus increases the probability that the same product will be selected in the future.

Johnson and Kotz (1977) show that the limit distribution of P_t for the Polya process is a Beta distribution with parameters $(b/s, w/s)$. With $b = w = s = 1$, then $P_t \sim \text{Beta}(1,1) = \text{Uniform}[0,1]$ as shown above. We now generalize our results based on the Polya process.

Theorem 3: The recommender does not influence the expected market shares of the products.

$E[P_t] = b/(b+w) = p$. Again, while the recommender does not influence the average outcome, realized outcomes can nonetheless be very different.

Modeling different market shares ($p \neq 1/2$): We can model a different value of p by changing the initial conditions (b, w) in the Polya model. Recall that $p = b/(b+w)$. In the Figure below, we plot the limiting distribution of the market shares (i.e., P_t) for three cases – $(b = 1, w = 1)$, $(b = 1, w = 2)$ and $(b = 1, w = 3)$ corresponding to $p=0.5$, $p=0.33$ and $p=0.25$ respectively. In case 1, the recommender has a homogeneity bias. In cases 2 and 3, there is a non-zero probability of the recommender reducing the market share differential as also a non-zero probability of a greater market share differential. Thus, in general, the same recommender is capable of either increasing or decreasing sales diversity. Further, there is strong path dependence in the realized outcome.

Deliberate burn-in time: The strong feedback effects observed are often driven by chance events early on. If firms are aware of this, a design choice at their disposal is allowing a burn-in period of data collection

to occur before the recommender is turned on. We can model this situation by holding s constant but increasing the initial values of (b, w) . Figure 4 demonstrates the impact of increasing the burn-in period. In accordance with theorem 3, recommenders do not influence the expected outcome. However, a longer burn-in period makes it less likely that the recommender will cause large deviations in the long-run market shares of the products. Nonetheless, recommenders continue to have the ability to either increase or decrease sales diversity.




Figure 3. Limiting distribution of market share with different initial conditions




Figure 4. Limiting distribution of market share for different burn-in periods

5 Conclusions

This paper studies the influence of recommender systems on sales diversity at online stores. Our results thus far show that for one common recommender system, sales diversity may actually decrease. However, the result is more subtle: while diversity could decrease on average for this recommender, there is a distribution over ultimate outcomes (diversity vs. homogeneity). Thus studies at individual firms can reveal either outcome, even though a bias may exist on average. For WISE 2006, we would like to model other recommender designs to see which designs lead to greater diversity. In addition, we are working to generalize our results and provide illustrations of the phenomenon. Specifically, we intend to generalize our model to a multi product case and explore ways to validate our model. These results would show that firms' often tacit design choices could have significant implications for patterns of product consumption.

Bibliography

- Anderson, C. 2004. The long tail. *Wired Magazine*.
- Arthur, W. B. 1990. Industry location and the importance of history. *Math Social Sciences* 19: 235-251.
- Breese, J., D., Heckerman, and C. Kadie. 1998. Empirical analysis of predictive algorithms for collaborative filtering. *14th Conference on Uncertainty in Artificial Intelligence*.
- Brynjolfsson, E., H. Yu, M. D. Smith. 2006. From Niches to Riches: The Anatomy of the Long Tail. *Sloan Management Review*, 47(4) 67-71.
- Cooke, A., H. Sujan, M. Sujan, and B. Weitz. 2002. Marketing the unfamiliar: the role of context and item-specific information in electronic agent recommendations. *Journal of Marketing Research* 39(4):488-497.
- Dorfman, R. 1979. A formula for the Gini coefficient. *The Review of Econ. and Statistics* 61(1):146-149.
- Johnson, N. L. and S. Kotz. 1977. *Urn Models and their Applications*. New York: John Wiley and Sons.
- R.J. Mooney and L. Roy. 1999. Content-based book recommending using learning for text categorization. In *SIGIR'99 Workshop on Recommender Systems: Algorithms and Evaluation*.
- Pham, A. and J. Healey. 2005. Tell you what you like. *Los Angeles Times*, 20 September.
- Schafer, J., J. Konstan, and J. Riedl. 1999. Recommender systems in e-commerce. In *Proceedings of the ACM Conference on Electronic Commerce*, p. 158-166.
- Senecal, S. and J. Nantel. 2004. The influence of online product recommendations on consumers' online choices. *Journal of Retailing* 80:159-169.