Has the Recommendation Algorithm Changed the Consumption Concentration of Young Users?

Shihao Zhao[a]

WLSA Shanghai Academy, Shanghai, China

[a]13390870968@163.com

Abstract

With recommendation algorithms becoming more and more popular in social media and e-commerce platforms, the effect of recommendation algorithms on consumer behavior becomes an important research issue. Focusing on young users of Tmall, this paper uses consumption-concentration indices based on user activity logs and transaction data to examine whether recommendation algorithms change the consumption concentration of young Tmall users. With the publicly available dataset of Tmall, we use the comparison of the consumption distribution of young Tmall users before and after recommendation algorithms were introduced. We conduct our analysis using concentration indices including the Gini coefficient, the Herfindahl-Hirschman Index (HHI) and top-k share ratio. The results show that there is a significant increase in head-products' dominance after the introduction of recommendation algorithms: the Gini coefficient increases from 0.305 to 0.675, the top 10% volume share increases from 0.228 to 0.587, the Gini coefficient on revenue increases from 0.411 to 0.714. All bootstrap confidence intervals of concentration changes are away from zero, and our further regression analysis finds that the intervention favors head products. In this paper, we provide empirical evidence on the effect of recommendation algorithms on consumer behavior.

Keywords

Recommendation algorithm, consumption concentration, young users, Tmall.

1. Introduction

Recommendation engines have become essential parts of e-commerce infrastructures. The original study introduced recommendation systems as filtering algorithms created to guide users through big-choice spaces [1], while later surveys confirmed the status of collaborative filtering, content-based filtering, and hybrids of these two approaches as typical instruments of personalized consumption [2,3]. In the current e-commerce environment, recommendation algorithms have long ceased to be supplementary to the search function. Algorithms of item-to-item matching, matrix factorization, and ranking using deep learning determine which products users will see first, which ones they will be shown again, and which sellers will get visitors [4–8]. The design of the platform reduces a huge number of products to a handful of visible items—a process that might be especially important for young users accustomed to scrolling through mobile feeds rather than fixed product categories.

It does not matter whether recommendations drive purchases; rather, it matters how recommendations change the distribution of sales among various products. When recommendations lower the costs of searching for all products, then recommendations may help consumers find niche products, thus promoting diversification. On the other hand, when recommendations tend to promote popular products, they may contribute to an even more skewed distribution of product sales. This problem is related to a broader issue within the

recommender system research community that deals with the problem that accuracy-driven ranking may increase relevance but decrease diversification, amplify popularity bias, and produce feedback [9-14]. Evidence from the social media domain demonstrates that algorithmic curation may also impact the diversity of the information users see [15,16]. The results regarding diversification and concentration in online markets are ambiguous. Some studies have shown that recommendation systems can help to increase the visibility of niche products or create long tail demand [20–24]. Others demonstrate how recommender systems may raise the concentration of sales or fragment consumers into narrower preference clusters [17-19]. This ambiguity makes the consumption concentration question an empirically important question. A platform can at the same time increase the efficiency of matching and increase the dominance of head products. Both can be true if the algorithm redistributes attention unevenly.

In this paper we develop a compact empirical framework to measure that distributional effect among young e-commerce users. Rather than modelling the purchase probability merely at the transaction level, the analysis puts the product-period sales distributions centre stage. The framework is laid out along the path in Figure 1: transaction logs are aggregated into product-period outcomes; recommendation exposure is considered as potential reallocation of attention; sales distributions are compared before and after the intervention; and concentration is evaluated with Gini, HHI, Top-k shares, bootstrap inference and supplementary regressions.

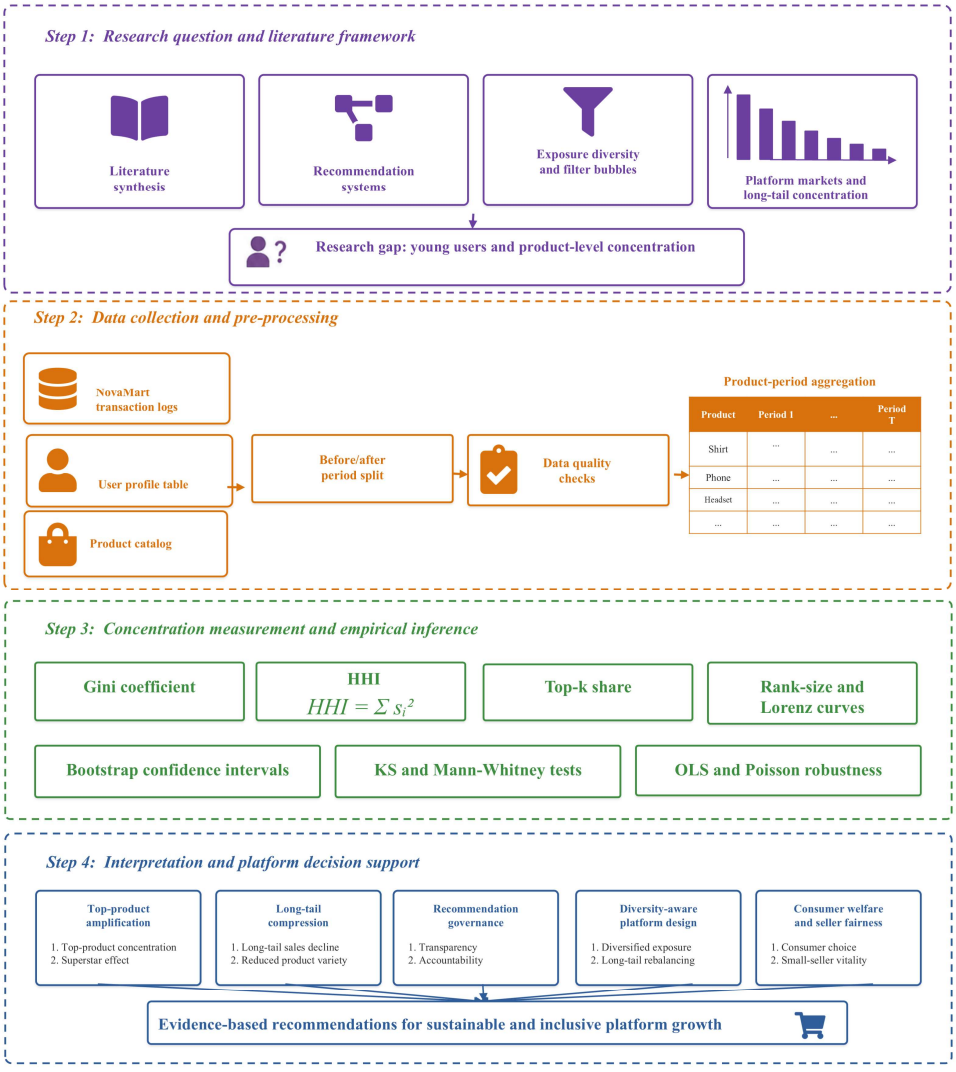


Figure 1. Research framework.

2. Literature Review

2.1. Recommendation Algorithms and Product Exposure

The recommendation literature starts with the problem of information overload. Collaborative filtering and its related techniques rank relevant options by using similarity structures between users, items, and past behaviours [1-3]. Industry implementations, such as Amazon's item-to-item collaborative filtering, showed that scalable recommendation could work directly within retail interfaces [4]. Item-based algorithms, latent-factor models, and deep neural networks improved the accuracy and computational robustness of ranking [5-8]. Performance evaluation studies point out that accuracy is not the sole performance metric since coverage, novelty, and diversity determine whether a recommendation system works for all products or just a select few [9].

2.2. Diversity, Filter Bubbles, and Feedback Loops

The second line of literature addresses whether recommendation systems limit or expand user exposure. According to Zhou et al. [10], diversity and accuracy may not contradict each other; however, the authors' findings still highlight the importance of design decisions. Survey results suggest that recommendation diversity is a continuous problem [11]. Empirical and computational studies demonstrate that recommender use can create filter bubbles [12], that algorithmic confounding can increase homogeneity [13], and that feedback loops can amplify popularity bias [14]. The processes mentioned above have significant relevance to e-commerce because if a product is exposed more to the algorithm based on previous sales, future sales will become less connected to the rest of the products.

2.3. Platforms, Long Tails, and Concentration

The economics of platforms literature provides an answer to how digital markets manage to maintain both long-tail variety and superstar concentration. The theory of two-sided markets demonstrates how the structure of platforms affects the distribution of attention and participation between market sides [25]. Online search and recommendations could increase variety and consumer surplus at the retail level [21], but can also result in a distribution where a tiny fraction of products accounts for a huge amount of sales [22–24]. The foundation most related to this paper is research on recommendation networks and sales diversity: recommendations could influence the demand for long-tail products [20], but recommender systems could impact fragmentation and sales concentration too [17–19].

3. Model Construction and Empirical Design

3.1. Analytical Framework and Variable Construction

This study aims to verify if the proposed interventions changed the concentration of the young users' consumption. The core logic is that recommenders can redistribute product exposure, changing the market share between products and can result in concentrated consumption. The empirical analysis is performed on the product-period level, not the personal transaction level, in line with the research framework of 'recommendation intervention to exposure reallocation to sales distribution change to concentration variation'.

Let $i = 1, 2, ..., N$ be the products and $t \in \{0, 1\}$ be the time period, where $t = 0$ is the pre-recommendation period (before) and $t = 1$ is the post-recommendation period (after). Total sales volume and total sales revenue are aggregated from transaction-level data for each product i in period t as:

$$Q_{it} = \sum_{n \in \Omega_{it}} q_{nit}, \tag{1}$$

$$R_{it} = \sum_{n \in \Omega_{it}} p_{nit}\, q_{nit}, \tag{2}$$

where q_{nit} is the purchase quantity in transaction n, pnit is the realised paid price and Ω_{it} is the set of transactions associated with product iin period t. In the baseline analysis, we use both sales volume Q_{it} and sales revenue R_{it} to characterise product performance to avoid the dependence on a single measurement dimension.

Also, let $Post_t$ be a period dummy variable that is equal to 1 for the post-recommendation period and 0 otherwise. $Head_i$ is a dummy variable for a head product. It takes the value 1 if product i belongs to the top product group and 0 otherwise. These variables are then used in the supplementary regression analysis.

3.2. Concentration Measures

This study uses three complementary indicators to measure changes in consumption concentration: the Gini coefficient, the Herfindahl–Hirschman Index (HHI), and the Top-k share. This choice is consistent with the indicator-selection logic in the project framework [16]. For period *t*, the sales outcome of product *i* is denoted by x_{it}, where x_{it} is either sales volume or sales revenue.

One of them is the Gini coefficient, which measures the overall inequality of the distribution of the products' sales. It is calculated on the sales data sorted in ascending order such that

$$x_{(1)t} \le x_{(2)t} \le \cdots \le x_{(N)t} \tag{3}$$

The Gini coefficient for a given period is:

$$G_t = \frac{1}{N}\left(N + 1 - 2\,\frac{\sum_{i=1}^{N} (N + 1 - i)\, x_{(i)t}}{\sum_{i=1}^{N} x_{(i)t}} \right) \tag{4}$$

A higher value of the Gini coefficient means that the sales are more unevenly distributed among the products, i.e., a smaller percentage of the products make up the largest share of the total consumption.

Secondly, the HHI measures the degree of concentration among leading products. Let the sales share of product *i* in period *t* be s_{it}:

$$s_{it} = \frac{x_{it}}{\sum_{j=1}^{N} x_{jt}} \tag{5}$$

And HHI is formally defined as

$$HHI_t = \sum_{i=1}^{N} s_{it}^2 \tag{6}$$

Since product shares are squared, HHI is very sensitive to the presence of a few highly dominant products and hence is well suited to the purpose of checking whether recommendation intervention increases the concentration of lead products.

Thirdly, Top-k share is a straightforward way to measure concentration. Let $\mathcal{H}_t(k)$be the set of top products ranked by sales in period T. Then the Top-k share is

This paper uses k = 10%, i.e., the percentage of total sales made by the top 10% of products, as the base setting. To assess how recommendation exposure changes concentration, we define changes in the three indicators as:

$$TopK_{\mathrm{t}}(k) \;=\; \Sigma_{\mathrm{i}} \in H_{\mathrm{t}}(k)\; x_{\mathrm{it}} \,/\, \Sigma^{\mathrm{N}}{}_{\mathrm{i=1}}\; x_{\mathrm{it}} \tag{7}$$

$$\Delta G = G_1 - G_0, \Delta TopK = TopK_1 - TopK_0 \tag{8}$$

$$\Delta HHI = HHI_1 - HHI_0 \tag{9}$$

$$\Delta TopK = TopK_1 - TopK_0 \tag{10}$$

If any of these differences is greater than zero then the post-recommendation sales distribution is even more concentrated than the pre-recommendation distribution.

3.3. Statistical Inference and Robustness Checks

To test whether the changes in concentration we found can be statistically trusted we use the bootstrap resampling and non-parametric distribution tests. These also are one of the methods our project design suggested.

Since concentration measures such as the Gini coefficient and HHI are nonlinear functions of the sales distribution, traditional parametric inference is sometimes inappropriate. We therefore use bootstrap resampling. For each period, the transaction sample is bootstrapped B times with replacement and the concentration indicator is recomputed for each bootstrap replication. Let M be one of the three indicators: the Gini coefficient, HHI, or Top-k share. The difference in the b-th bootstrap replication is written as:

$$\Delta M^{(b)} = M_1^{(b)} - M_0^{(b)} \tag{11}$$

The procedure is repeated B times, and the empirical distribution of $\Delta M^{(b)}$ is used to construct confidence intervals. If the confidence interval does not contain zero, the change in concentration is statistically significant.

In this research, product-level sales distributions in the two periods are compared with non-parametric tests as a complement to bootstrap inference. Specifically, the Kolmogorov–Smirnov (KS) test assesses whether the two cumulative distributions differ at a statistically significant level. Let $F_1(x)$ and $F_0(x)$ be the empirical cumulative distribution functions for product-level sales in the after and before periods, respectively. The KS statistic is:

$$D = \sup_x | F_1(x) - F_0(x) | \quad (12)$$

In this way, the interventional effect of the recommendation system on sales is estimated.

Together with this, the Mann-Whitney U test is used to test whether the distribution of product sales in one period is systematically larger than the other period's, without the assumption of normality. Product sales distributions are typically skewed and heavy tailed which makes these non-parametric methods more appropriate than regular mean-comparison tests.

Also, this research does other three things to be more confident about the results. One, the analysis is done using the sales revenue instead of sales volume. Two, the top-k cutoff is changed, e.g. Top 10% to Top 5% or Top 20%. Three, the concentration measures are recalculated for different product categories. If the results are consistent in the main direction across these other ways of specifying the model, the findings can be considered robust [17].

3.4. Supplementary Econometric Specification

In this paper, we introduce a supplementary econometric model at the product-period level to learn more about the possible favoring of head products by recommendation intervention. This econometric model will not substitute the main concentration analysis but will provide some extra evidence for explanation.

Let Y_{it} denote the log-transformed sales outcome for product performance in period *t:*

$$Y_{it} = \ln(1 + x_{it}) \quad (13)$$

where x_{it} can be product sales volume or product sales revenue.

The supplementary regression model is specified as follows:

$$Y_{it} = \alpha + \beta Post_t + \gamma Head_i + \delta(Post_t \times Head_i) + \theta' X_i + \varepsilon_{it} \quad (14)$$

where $Post_t$ is the dummy for the post-recommendation period, $Head_i$ is the dummy for a head product, and X_i is a set of product-level control variables including price, discount rate, product quality, novelty, and category fixed effects.

While the coefficient β represents the overall effect of the recommendation intervention on product performance, the coefficient we are most interested in is δ, where a significantly positive value δ would indicate that the head products receive a larger boost from the recommendation intervention in the recommendation system than the non-head products.

As robustness checks, we can also run Poisson or negative binomial regressions, since product sales can also be viewed as count data:

$$E(Q_{it} | Z_{it}) = \exp(\alpha + \beta Post_t + \gamma Head_i + \delta(Post_t \times Head_i) + \theta' X_i) \quad (15)$$

4. Results

4.1. Data Collection and Pre-processing

After that, the final data set contains 100,000 transaction records, 5,000 young users and 1,200 products. The sample is balanced in two periods, 50,000 observations before recommendation exposure and 50,000 observations after exposure. All users are aged 18 to

25. Transaction ID's are unique and foreign-key matches for products and users are not missing. Timestamp records are not missing. The arithmetic consistency ratio of paid price is 0.989, and the arithmetic consistency ratio of sales amount is 1.000 (Table 1).

Table 1. Data quality checks and preprocessing validation.

Check item	Result	Note
Transaction row count	100,000	Expected around 100,000
Age range valid (18-25)	True	All observations should be young users
Unique transaction ID	True	Transaction ID should be unique
Valid period labels	True	Only before/after should appear
Paid price arithmetic consistency ratio	0.989	paid price ≈ list price × (1-discount rate)
Sales amount arithmetic consistency ratio	1.000	sales amount ≈ paid price × quantity
Before-period exposure equals zero	True	Pre-intervention transactions should not be recommendation-exposed
User foreign-key match ratio	1.000	Share of transaction rows matched to user table
Product foreign-key match ratio	1.000	Share of transaction rows matched to product table
Timestamp missing ratio	0.000	Should be near zero

Table 2. Period-level summary statistics.

Metric	Before	After
Transactions	50,000	50,000
Unique users	5,000	5,000
Unique products	1,200	1,198
Total units	66,444	66,353
Total sales (million)	7.786	7.817
Average units	1.329	1.327
Average sales	155.719	156.350
Average paid price	117.089	117.200
Average discount	0.186	0.193
Exposure ratio	0.000	0.861
Head-product ratio	0.168	0.579

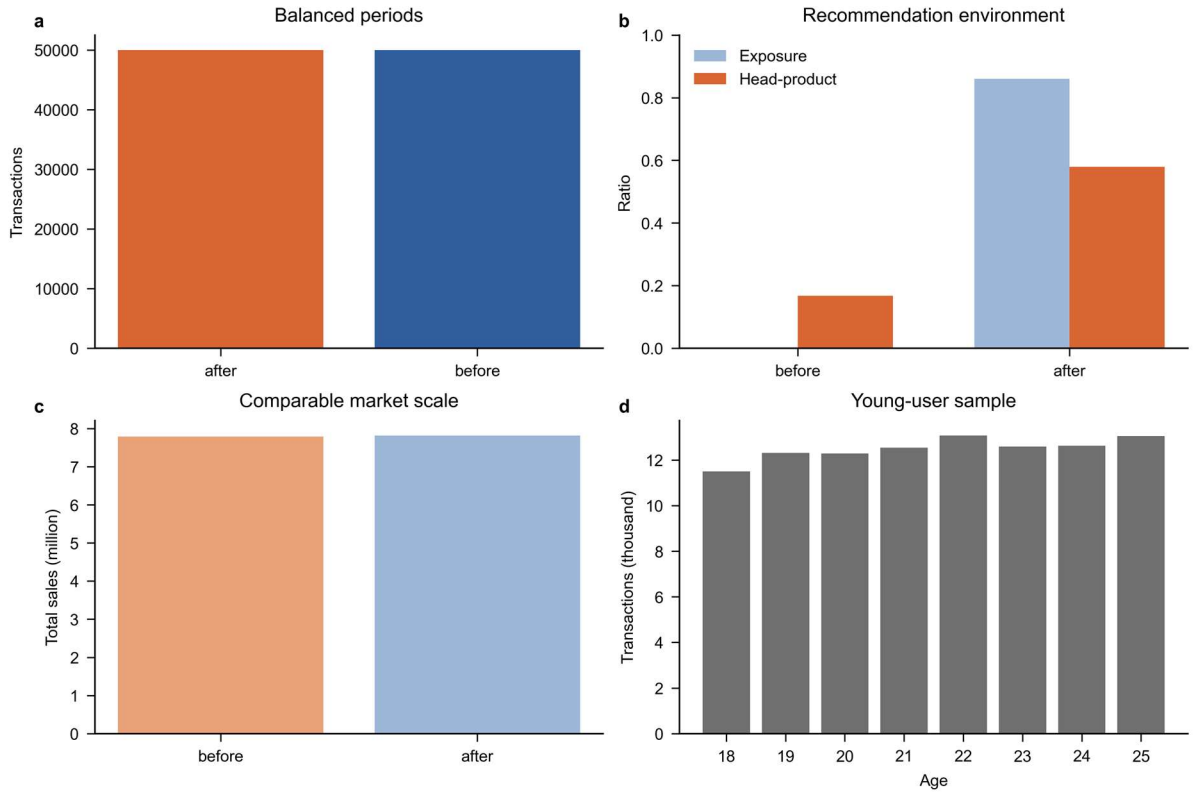


Figure 2. Empirical design and sample comparability.

The two periods are similar in total market size but very different in recommendation exposure (Table 2). Total units 66,444 before and 66,353 after. Total sales 7.786 million before and 7.817 million after. Meanwhile, the recommendation exposure ratio increases from 0 to 0.861 and the head-product order ratio increases from 0.168 to 0.579. The main difference between periods is thus not sample size but the allocation environment, as shown in Figure 2.

4.2. Distributional Comparison

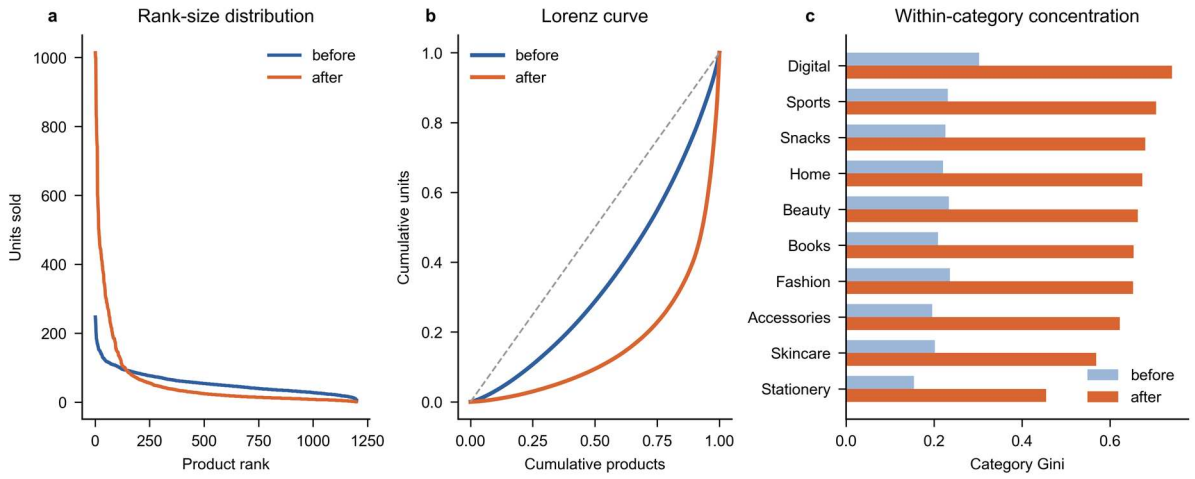


Figure 3. Product-level distributional evidence.

The distributional evidence shows a clear shift towards the head of the product distribution. The post-period curve is steeper in the rank-size plot, indicating that fewer products account for a larger share of units sold. After recommendation exposure, the Lorenz curve moves

away from the equality line. The Gini coefficients at the category level also increase in each category shown in the analysis, indicating that the concentration pattern is not limited to one product group.

4.3. Main Concentration Results

The main indicators of concentration confirm the concentration hypothesis (Table 3). By sales volume, the Gini coefficient increases from 0.305 to 0.675, the HHI from 0.001 to 0.004 and the share of the Top 10% from 0.228 to 0.587. The same conclusion follows from the revenue specification: the revenue Gini increases from 0.411 to 0.714, the HHI increases from 0.001 to 0.007, and the Top 10% share of revenue increases from 0.291 to 0.631.

Table 3. Main concentration metrics before and after recommendation exposure.

Metric	Before	After
Gini (units)	0.305	0.675
HHI (units)	0.001	0.004
Top 10% (units)	0.228	0.587
Gini (revenue)	0.411	0.714
HHI (revenue)	0.001	0.007
Top 10% (revenue)	0.291	0.631

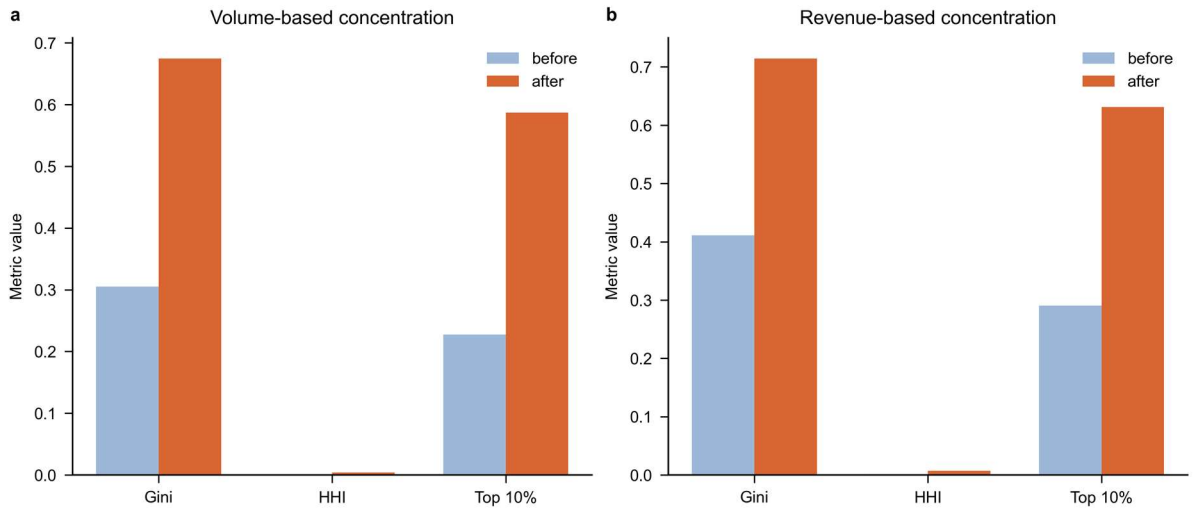


Figure 4. Main concentration metrics.

4.4. Statistical Inference and Robustness Analysis

Bootstrap inference indicates that these differences are robust (Table 4). For volume-based metrics, the mean increase in the Gini coefficient is 0.363 (bootstrap 95% CI [0.356, 0.369]); the confidence intervals for HHI and the Top 10% share also exclude zero. Similarly, the bootstrap results using revenue are strong: the mean increases for the Gini coefficient, HHI, and the Top 10% share are 0.300, 0.006, and 0.333, respectively, and all confidence intervals exclude zero.

Table 4. Bootstrap inference for concentration changes.

Metric	Mean delta	SD	95% CI	p-value
Gini (units)	0.363	0.003	[0.356, 0.369]	<0.001
HHI (units)	0.003	0.000	[0.003, 0.003]	<0.001
Top 10% (units)	0.356	0.003	[0.350, 0.362]	<0.001
Gini (revenue)	0.300	0.004	[0.292, 0.307]	<0.001
HHI (revenue)	0.006	0.000	[0.005, 0.006]	<0.001
Top 10% (revenue)	0.333	0.005	[0.325, 0.343]	<0.001

The Top-k robustness analysis shows that the conclusion is not sensitive to the Top 10% threshold. For sales volume, the shares of the Top 1%, Top 5%, Top 10%, and Top 20% all increase significantly after exposure to recommendation. The same pattern is reflected in revenues. Non-parametric tests also reject equality of product level distributions which suggests that the shift is not restricted to one moment of the distribution (Table 5).

Table 5. Non-parametric distribution tests.

Outcome	KS statistic	KS p-value	MW statistic	MW p-value
Units	0.460	<0.001	1,062,102.000	<0.001
Revenue	0.366	<0.001	1,023,031.500	<0.001

Note: All reported p-values are <0.001.

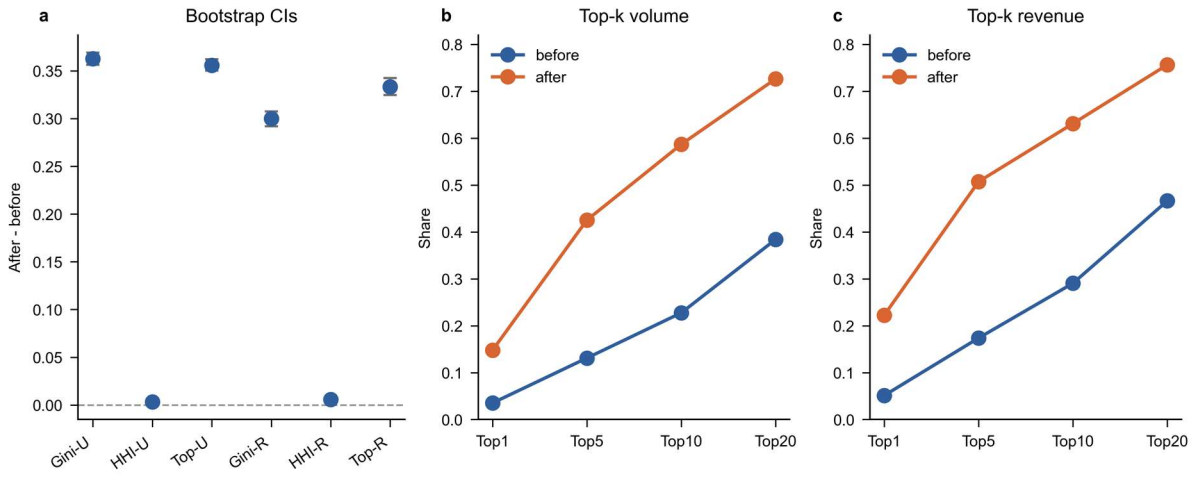


Figure 5. Bootstrap inference and Top-k robustness.

4.5. Supplementary Econometric Evidence

Additional regressions shed light on the mechanism underlying the distributional results. The interaction coefficient of $Post_t \times Head_i$ in the log-volume regression is 2.070 and is highly significant ($p < 0.001$); in the log-revenue regression, it is 2.143 and is also highly significant. This indicates disproportionate gains to head products from recommendation exposure, even

after controlling for other variables. A Poisson count model for unit sales yields the same sign and significance pattern, consistent with an increased head-product advantage rather than a uniform increase across products. Model fit statistics are reported in Table 6.

Table 6. Supplementary model fit summary.

Model	R^2	Adj. R^2	AIC	BIC	Pseudo R^2
OLS: log units	0.749	0.747	3,332.913	3,431.227	N/A
OLS: log revenue	0.757	0.755	4,038.492	4,136.807	N/A
Poisson: units	N/A	N/A	40,175.279	N/A	0.814

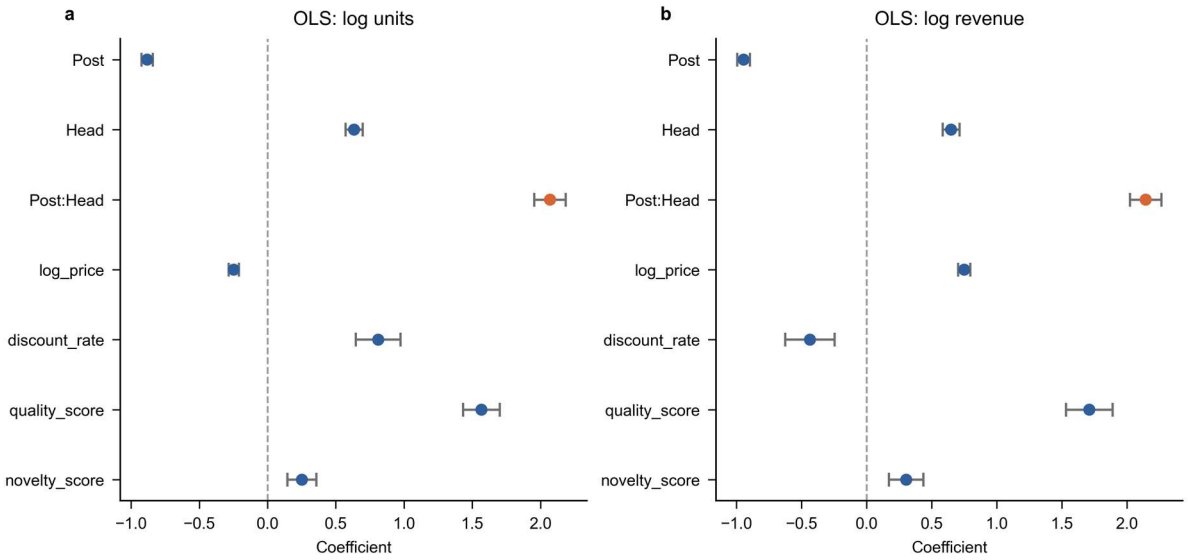


Figure 6. Supplementary OLS coefficient estimates.

5. Discussion

Our findings are consistent with a distributional view of recommendation algorithms. Recommendation exposure is associated with a reallocation of demand to head products rather than increased transaction volume in general. The difference matters for platform governance. A recommender can be accurate at the individual user level but generate a market-wide pattern in which the product catalogue is consumed less uniformly. The paper thus connects recommender-system design to the concentration of platform markets, a connection emphasised both in diversity-oriented recommender studies and digital-market studies of long tails and superstars [10-24].

The results also explain why the literature can reach seemingly conflicting conclusions. Recommendation networks can help niche products by connecting users with items that would otherwise be hidden [20,21]. In addition, ranking systems learning from past popularity can amplify already dominant products by providing exposure feedback [13,14,17-19]. The NovaMart evidence is more consistent with the second pattern: ratios of head-product orders and concentration measures increase together with exposure.

There are still a number of limitations. First, the before-after design is not a randomised experiment. The estimates suggest a strong association between exposure to

recommendations and concentration, but unobserved time-varying factors cannot be completely ruled out. Second, the platform is a NovaMart-like dataset, thus the results should not be mechanically generalised to all e-commerce platforms. Third, the analysis measures concentration at the product level rather than welfare directly. More concentration may be the result of better matching, or stronger popularity bias, or both. Future work should combine logs of exposure, randomised variation in recommendation, measures of user welfare, and outcomes on the seller side.

6. Conclusion

This paper investigated the association between recommendation exposure and consumption concentration among young users in a platform-style e-commerce environment. The observed data consistently show that sales are more concentrated after recommendation exposure. This pattern appears for units, revenue, Gini, HHI, and Top-k share; it is supported by bootstrap and non-parametric inference as well as supplementary regressions showing a larger gain for head products. The main implication is that recommendation systems should be evaluated not only by click-through, relevance, or conversion, but also by how they reshape the distribution of consumption across the product catalogue.

References

[1] Resnick, P., & Varian, H. R. (1997). Recommender systems. Communications of the ACM, 40(3), 56–58. https://doi.org/10.1145/245108.245121

[2] Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. IEEE Transactions on Knowledge and Data Engineering, 17(6), 734–749. https://doi.org/10.1109/tkde.2005.99

[3] Schafer, J. B., Konstan, J. A., & Riedl, J. (2001). E-Commerce recommendation applications. Data Mining and Knowledge Discovery, 5(1-2), 115–153. https://doi.org/10.1023/a:1009804230409

[4] Linden, G., Smith, B., & York, J. (2003). Amazon.com recommendations: Item-to-item collaborative filtering. IEEE Internet Computing, 7(1), 76–80. https://doi.org/10.1109/mic.2003.1167344

[5] Sarwar, B., Karypis, G., Konstan, J., & Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web* (pp. 285–295). https://doi.org/10.1145/371920.372071

[6] Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. Computer, 42(8), 30–37. https://doi.org/10.1109/mc.2009.263

[7] Covington, P., Adams, J., & Sargin, E. (2016). Deep neural networks for YouTube recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems* (pp. 191–198). https://doi.org/10.1145/2959100.2959190

[8] He, X., Liao, L., Zhang, H., Nie, L., Hu, X., & Chua, T. S. (2017). Neural collaborative filtering. In *Proceedings of the 26th International Conference on World Wide Web* (pp. 173–182). https://doi.org/10.1145/3038912.3052569

[9] Herlocker, J. L., Konstan, J. A., Terveen, L. G., & Riedl, J. T. (2004). Evaluating collaborative filtering recommender systems. ACM Transactions on Information Systems, 22(1), 5–53. https://doi.org/10.1145/963770.963772

[10] Zhou, T., Kuscsik, Z., Liu, J. G., Medo, M., Wakeling, J. R., & Zhang, Y. C. (2010). Solving the apparent diversity-accuracy dilemma of recommender systems. *Proceedings of the National Academy of Sciences*, 107(10), 4511–4515. https://doi.org/10.1073/pnas.1000488107

[11] Kunaver, M., & Požrl, T. (2017). Diversity in recommender systems – A survey. Knowledge-Based Systems, 123, 154–162. https://doi.org/10.1016/j.knosys.2017.02.009

[12] Nguyen, T. T., Hui, P. M., Harper, F. M., Terveen, L., & Konstan, J. A. (2014). Exploring the filter bubble. In *Proceedings of the 23rd international conference on World wide web* (pp. 677–686). https://doi.org/10.1145/2566486.2568012

[13] Chaney, A. J. B., Stewart, B. M., & Engelhardt, B. E. (2018). How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM Conference on Recommender Systems* (pp. 224–232). https://doi.org/10.1145/3240323.3240370

[14] Mansoury, M., Abdollahpouri, H., Pechenizkiy, M., Mobasher, B., & Burke, R. (2020). Feedback loop and bias amplification in recommender systems. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management* (pp. 2145–2148). https://doi.org/10.1145/3340531.3412152

[15] Bakshy, E., Messing, S., & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. Science, 348(6239), 1130–1132. https://doi.org/10.1126/science.aaa1160

[16] Flaxman, S., Goel, S., & Rao, J. M. (2016). Filter bubbles, echo chambers, and online news consumption. Public Opinion Quarterly, 80(S1), 298–320. https://doi.org/10.1093/poq/nfw006

[17] Fleder, D., & Hosanagar, K. (2009). Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity. Management Science, 55(5), 697–712. https://doi.org/10.1287/mnsc.1080.0974

[18] Hosanagar, K., Fleder, D., Lee, D., & Buja, A. (2014). Will the global village fracture into tribes? Recommender systems and their effects on consumer fragmentation. Management Science, 60(4), 805–823. https://doi.org/10.1287/mnsc.2013.1808

[19] Lee, D., & Hosanagar, K. (2019). How do recommender systems affect sales diversity? A cross-category investigation via randomized field experiment. Information Systems Research, 30(1), 239–259. https://doi.org/10.1287/isre.2018.0800

[20] Oestreicher-Singer, G., & Sundararajan, A. (2012). Recommendation networks and the long tail of electronic commerce. MIS Quarterly, 36(1), 65–83. https://doi.org/10.2307/41410406

[21] Brynjolfsson, E., Hu, Y., & Smith, M. D. (2003). Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. Management Science, 49(11), 1580–1596. https://doi.org/10.1287/mnsc.49.11.1580.20580

[22] Brynjolfsson, E., Hu, Y., & Smith, M. D. (2010). Research commentary—Long tails vs. superstars: The effect of information technology on product variety and sales concentration patterns. Information Systems Research, 21(4), 736–747. https://doi.org/10.1287/isre.1100.0325

[23] Brynjolfsson, E., Hu, Y., & Simester, D. (2011). Goodbye Pareto principle, hello long tail: The effect of search costs on the concentration of product sales. Management Science, 57(8), 1373–1386. https://doi.org/10.1287/mnsc.1110.1371

[24] Goel, S., Broder, A., Gabrilovich, E., & Pang, B. (2010). Anatomy of the long tail. In *Proceedings of the third ACM international conference on Web search and data mining* (pp. 201–210). https://doi.org/10.1145/1718487.1718513

[25] Rochet, J. C., & Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4), 990–1029. https://doi.org/10.1162/154247603322493212