WEP
Warwick
Evans
Publishing

Transactions on Economics, Business and Management Research
ISSN: 2960-1789, eISSN: 2960-2254 | Volume 17
WEPM 2026

Does Algorithmic Recommendation Weaken Consumers' Original Preferences?

-- An Economic Analysis Based on Information Constraints

Qiyue Zhang

School of Art and Science, University of North Carolina Chapel Hill, Chapel Hill, 27514, USA

Abstract. The paper explores the topic of whether algorithmic recommendation systems radically change the preferences of consumers by constraining information. We construct a theoretical model of search costs, dynamic learning and endogenous preference formation, in which algorithms act as information intermediaries that selectively filter information about available options in a systematic manner. On the basis of extensive e-commerce purchases records of 35,874 purchases made by 1,000 participants over a period of 12 months, we apply difference-in-differences techniques to determine causal impacts. Our findings indicate that algorithmic exposure enhances concentration of consumption by 52% (measured using Herfindahl Index) and decreases the diversity of preferences by 91 percent (measured using Shannon entropy) and both are statistically significant at p<0.001. The analysis of mechanisms shows that feedback loops have 45% of the total effect, the information filtering has 35%, and cognitive dependence characterizes 20%. The effect of treatment has a heterogeneous search cost sensitivity and accrual over time due to the accumulation of user history over time among algorithms. We do not make any evidence of other readings like refinement of tastes or supply limitation. The results of this research contradict the traditional perspective that algorithms only support the expression of preference, pointing rather to its active involvement in the formation of preferences with unclear welfare implications. Some policy recommendations suggested are platform design that promotes diversity, compulsory impact evaluation of algorithms, and consumer empowerment interventions.

Keywords: Algorithmic recommendations, preference formation, filter bubbles, consumer behavior, information constraints.

1. Introduction

The emergence of algorithmic recommendation systems has radically altered the consumer experience of product, service, and content discovery and interaction in digital platforms. In e-commerce giants such as Amazon or streaming services such as Netflix to social media platforms such as Tik Tok, recommendations algorithms have become ubiquitous between customers and massive catalogs of options. These systems are using advanced machine learning methods such as collaborative filtering, content-based filtering and hybrid methods to predict and recommend item that are relevant to the historical preferences and behaviors of users. Although this kind of personalization has the potential to provide greater user satisfaction and lower search expenditures, a question that arises is, do these algorithms only help in the expression of preferences or in fact they are changing the very nature of consumer preferences themselves?

Empirical evidence recently has indicated that algorithmic recommendations can cause significant and potentially distortive effects on consumer preference formation. As shown in agent-based simulations by Anwar et al. [1], the traditional recommendation algorithms may have severe effects on the inter-user and intra-user consumption diversity, with the trade-offs between the effects of filter bubbles and content homogenization being more intricate. Their results show that behavior-based recommendations made in the past decrease the filter bubbles by influencing diversity among users and not necessarily among the individual patterns of consumption. On the same note, a study of consumer behavior in AI-based systems demonstrates that customized recommendation systems have the potential to generate so-called filter bubbles in which consumers are systematically isolated within their informational comfort zones, resulting in limited exposure to alternative options [2]. The phenomenon brings up some underlying issues on consumer welfare: in the event that algorithms cause consumers to reinforce their existing preferences at the expense of exploration, there is a risk that consumers will become confined in a progressively smaller consumption channel, and some long-term utility might be traded-off with short-term satisfaction.

Economically, this is a break with classical types of consumer choice which are held to assume stable, well-defined consumer preferences generated outside of the context of choice. The conventional neoclassical model is the one in which the preferences are exogenous i.e. they are fixed tastes that direct the consumption decision within resource limitations. Nonetheless, the interactiveness of algorithmic recommendation systems opens a process of endogenous preference formation, in which repeated exposure to the algorithm can systematically change consumer desired preferences over time. This interaction poses both theoretical and empirical difficulties: we need to construct frameworks with the ability to accommodate the evolution of preferences when we are constrained by information; and we need to estimate and measure these effects with the help of proper methodological instruments. Most recent research has investigated the impact of algorithmic affordances on preventing the creation of filter bubbles [3] proposes that the design of systems should be careful in order to prevent certain adverse effects, whereas the research done on impulse purchasing behavior demonstrates that the quality of recommendations and their diversity directly impact consumer decision-making behavior [4].

This paper constructs an extensive economic model to determine the impact of algorithmic recommendation systems on the consumer preference based on information limitations. We contribute in three main aspects: First, we develop a theoretical framework of the relationship between algorithmic information filtering, search costs, and preference evolution, which generalizes the traditional consumer choice theory to a dynamic process of preference formation. Second, we offer empirical testing of major predictions of our model with real-world data on e-commerce to test the effects of recommendation exposure on preference diversity by using difference-in-differences and regression discontinuity methods. Third, we examine the precise ways in which algorithms shape preferences information filtering at the cost of search, feedback mechanisms that support already made decisions and cognitive reliance on algorithmic defaults and explain their effects on consumer welfare and platform regulation. We propose that, although algorithmic recommendations can be effective in the short-term, they can create significant long-term expenses of narrowing preferences and decreased consumer autonomy, and this practice should be considered in platform design and policy-making.

2. Literarure Review

The current recommendation systems utilize three basic strategies, namely, collaborative filtering, content-based filtering and a mix of both. Collaborative filtering examines user trends to provide predictions and find similarities whereas content-based methods use item characteristics and user profiles. Ali et al. [5] offer an extensive review that shows machine learning and deep learning improvements have significantly boosted the performance of collaborative filtering, especially with regard to conventional problems like sparsity of data and cold-start issues. New developments combine neural collaborative filtering with recurrent networks to learn non-linear relationships and sequential patterns of behavior in a complex way [6]. Such hybrid architectures have been particularly useful in large-scale tasks and have delivered quantifiable performance advantages in rating accuracy and ranking work on a wide range of datasets.

One important issue that has arisen as a result of the literature on algorithmic recommendations is the establishment of filter bubbles, where personalization algorithms selectively depersonalize users into different content. M. G. Younis [7] overview this phenomenon in a systematic fashion, where all three issues mentioned (reduced content diversity, privacy issues, and bias of the algorithm) are reported. Their in-depth assessment establishes the filter bubbles as a general design issue that needs to be addressed to fairness, transparency, and trust. The mechanism works in a feedback loop: by emphasizing the content that is similar to previous interactions, the user consumes recommended content, and the recommendation becomes even more homogeneous. This cycle can significantly reduce the decisions available to consumers making it possible to limit the exploration of preferences and identification of new alternative options.

The classical theory of consumer choice holds the assumption that preferences are exogenous i.e. fixed and constant across situations of decision. Behavioral economics however does not acknowledge this presumption by showing that preferences tend to be endogenous, or affected by contexts, presentation of information, and previous experiences. P. J. Bus. [8] reviews experimental data indicating that consumers are systematic deviators with cognitive biases, framing effects and social influences. The endogeneity of preferences suggests that the economic, social and technological systems have the potential to create and determine the consumption tastes and behaviors. M. D. [9] critically evaluates this difficulty to the welfare analysis by putting forward the idea that the process of forming their preferences is complicated and context specific and tends to be made through subtle mechanisms that decision-makers are unaware of. Within the framework of algorithmic systems, this speculative viewpoint implies that recommendation algorithms could be not only helpful towards constructing pre-existing preferences but also involved in their formation and development through time.

Even though there has been significant advancements in the understanding of the algorithmic recommendation systems, and the study of behavioral economics separately, there is limited literature in systematic economic models that examine their intersection. These two aspects are largely covered in the current literature: either technical performance optimization [10] or qualitative description of filter bubble effects, but there are no formal models that would connect algorithmic information constraints with preference evolution. Moreover, there is a lack of empirical methods of identification that can plausibly confirm the existence of causal links between recommendation exposure and preference change. This paper addresses these gaps by developing a tractable theoretical model incorporating information constraints, search costs, and dynamic preference formation, complemented by rigorous empirical analysis using modern identification techniques to quantify the magnitude and mechanisms of algorithmic influence on consumer preferences.

3. Theoretical Framework and Model

3.1 Basic Setup and Utility Function

We model a consumer who makes sequential consumption choices over a finite time horizon $t \in \{0, 1, ..., T\}$. At each period, the consumer chooses from a set of *alternatives* $X = \{x_1, x_2, ..., x_n\}$, where each alternative x_i is characterized by a vector of attributes. The consumer's underlying preferences are represented by a utility function:

$$U(x_i, \theta) = \alpha \cdot v(x_i) + \beta \cdot f(x_i, \theta) + \varepsilon_i \quad (1)$$

Where $v(x_i)$ represents the intrinsic value of alternative i, $\theta \in \Theta$ is the consumer's type parameter capturing individual heterogeneity in preferences, $f(x_i, \theta)$ is an idiosyncratic preference component that varies with consumer type, and ε_i is an i.i.d. random utility shock following a *Type I* extreme value distribution. The parameters α and β determine the relative weights of intrinsic value and type-specific preferences. In the absence of information constraints, the consumer would evaluate all alternatives and select the one maximizing (1).

3.2 Information Constraints and Search Costs

Evaluating all alternatives is costly. We denote the consumer's information set at time t as $S_t \subseteq X$, representing the subset of alternatives the consumer has examined. Acquiring information about alternative i involves a search cost $c(i) > 0$. Without algorithmic assistance, the consumer faces a sequential search problem: in each period, she decides whether to continue searching (expanding S_t) or to select from the currently known alternatives. The optimal stopping rule balances the marginal benefit of discovering potentially superior alternatives against search costs.

Let $V(S_t)$ denote the value function associated with information set S_t:

$$V(S_t) = max\{max_{i \in S_t} U(x_i, \theta), E[V(S_{t+1})|S_t] - \bar{c}\} \tag{2}$$

Where $\bar{c}$ represents the expected cost of expanding the information set. The consumer stops searching when the expected value of continuing search falls below the utility of the best currently known alternative.

3.3 Algorithmic Recommendation as Information Intermediary

We formalize the recommendation algorithm as an information intermediary that reduces search costs but introduces systematic filtering. The algorithm observes the consumer's history $H_t = \{x_0, x_1, ..., x_{t-1}\}$ of past choices and generates a recommended set $R_t \subseteq X$. The algorithm's objective is to maximize a platform utility function:

$$\Pi_a(R_t|H_t) = E[U(x_i, \theta)|H_t, x_i \in R_t] + \lambda \cdot P(engagement|R_t) \tag{3}$$

Where λ determines the weight on platform-specific objectives such as engagement or click-through rates. The algorithm uses collaborative filtering to identify patterns: for consumer with history H_t, it recommends items consumed by similar users or items similar to past choices. Mathematically, similarity is measured through a distance metric $d(x_i, x_j)$ in attribute space, and the algorithm recommends items minimizing $\Sigma_{j \in H_t} d(x_i, x_j)$.

Crucially, algorithmic recommendations reduce effective search costs for items in R_t while increasing costs for items outside R_t. We model this as:

$$c(i|R_t) = \{c_L \; if \; i \in R_t; c_H \; if \; i \notin R_t\} \tag{4}$$

Where $c_L < \bar{c} < c_H$. This formalization captures the dual effect of recommendations: they facilitate discovery within a narrow subset while making alternatives outside this subset effectively invisible.

3.4 Filter Bubble Effect and Preference Narrowing

The filter bubble emerges from the interaction between algorithmic filtering and consumer search behavior. We define the consumer's effective choice set at time t as $S_t^{eff} = S_t \cap R_t$, representing alternatives both known and recommended. Over time, repeated exposure to algorithmically filtered content causes S_t^{eff} to concentrate around past choices.

We measure preference diversity using the Herfindahl-Hirschman Index (HHI) of consumption shares across product categories:

$$HHI_t = \Sigma_k s_{kt}^2 \tag{5}$$

Where s_{kt} is the share of consumption in category k at time t. Higher HHI indicates greater concentration and reduced diversity. We predict that algorithmic exposure increases HHI over time as consumers concentrate purchases within narrower categories.

Alternatively, we employ Shannon entropy as a diversity measure:

$$H_t = -\Sigma_k s_{kt} \cdot log(s_{kt}) \tag{6}$$

Lower entropy indicates reduced diversity. Both metrics capture the key prediction: algorithmic recommendations systematically narrow the distribution of consumption over time.

3.5 Dynamic Preference Formation and Bayesian Updating

Preferences evolve endogenously through Bayesian learning. The consumer is initially uncertain about her true type θ and updates beliefs based on experienced utility from consumption. Let μ_t represent the consumer's belief about θ at time t, with prior distribution $\mu_0 \sim N(\theta_0, \sigma_0^2)$. After consuming x_i and observing utility u_i, the consumer updates beliefs according to Bayes' rule:

$$\mu_{t+1} = (\tau_t\mu_t + \tau_s u_i)/(\tau_t + \tau_s) \tag{7}$$

Where $\tau_t = 1/\sigma_t^2$ is the precision of prior beliefs and τ_s is the precision of the utility signal. Posterior beliefs μ_{t+1} are a precision-weighted average of prior beliefs and observed outcomes [11]. Critically, if algorithmic recommendations limit the consumer to a narrow subset of alternatives, learning occurs only within this restricted domain. The consumer's beliefs about θ become increasingly specialized to predict utility within R_t but provide little guidance for alternatives outside this set. This creates path dependence: early algorithmic exposure shapes the direction of belief updating, which subsequently reinforces the tendency to consume similar items.

3.6 Equilibrium Characterization

We characterize the steady-state equilibrium where the consumer's beliefs, information set, and algorithm recommendations stabilize. In equilibrium:

1). The consumer optimally chooses from her effective choice set: $x^*_t \in argmax_{x \in S_t^{eff}} E[U(x, \theta)|\mu_t]$

2). The algorithm generates recommendations that maximize (3) given observed history

3). Consumer beliefs stabilize: $\mu_{t+1} \approx \mu_t$ (convergence in posterior means)

The model yields several testable predictions. First, exposure to algorithmic recommendations should increase consumption concentration (higher HHI, lower entropy) relative to consumers with unrestricted search. Second, the effect should be stronger for consumers with higher search costs or lower prior knowledge, as these individuals rely more heavily on algorithmic guidance. Third, diversity should decline more rapidly in the early periods of algorithm exposure, with the rate of decline slowing as beliefs and consumption patterns stabilize.

3.7 Welfare Implications

The benefits of welfare are not clear. Search costs and utility are decreased and high-utility alternatives are discovered in the short run due to algorithms. It is however the subject of the long-run welfare that the true optimal bundle of the consumer falls within or outside the algorithmic filter. When algorithmically the better alternatives are systematically left out, the consumer will suffer a continuing welfare loss that will accumulate with time as beliefs become more and more off-optimal. We model this trade-off by comparing the expected lifetime utility with algorithmic exposure relative to unrestricted search taking into consideration both the short-term returns associated with lower search costs and the long-term losses associated with preference distortion.

4. Data, Methodology, and Results

4.1 Data Sources and Description

The current empirical study applies a rich set of e-commerce transactions that are aimed at simulating the organization and features of the real online shopping environment. The sample includes 35,874 purchase records of 1,000 customers in 500 products of 20 different product categories in a 12-month follow up. The users will be assigned randomly to either a treatment group (n=700) that is subjected to the algorithmic recommendations or the control group (n=300) that walks without algorithmic support. This quasi-experiment design will allow establishing the causal relationship between the effects of algorithms on preference formation.

Realistic behavioral heterogeneity is included to the data generation process. Latent characteristics of each user include the search cost sensitivity (based on Beta (2, 5) distribution, which includes different degrees of readiness to look at alternatives) and exploration tendency (based on Beta (2,2) distribution, which is an inquisitive attitude to new product categories). In the treatment group, individual recommendations are made based on collaborative filtering algorithms that determine the patterns of the similar users and suggest the items with attribute similarity to the past purchases.

Notably, the algorithm parameters are adjusted to depict the real e-commercerecommendation systems reported in the industry literature, which guarantees the external validity of our empirical design.

4.2 Empirical Strategy

Our primary identification strategy employs a difference-in-differences (DID) framework that exploits temporal variation in algorithmic exposure intensity combined with cross-sectional variation in treatment assignment. The baseline specification is:

$$Diversity_{it} = \alpha + \beta_1 \cdot Treatment_i + \beta_2 \cdot Post_t + \beta_3 \cdot (Treatment_i \times Post_t) + \gamma \cdot X_{it} + \varepsilon_{it} \quad (8)$$

Where Diversity_{it} represents preference diversity for user *i* at time *t* (measured via HHI or Shannon entropy), Treatment_i is an indicator for algorithmic exposure, Post_t indicates periods after the midpoint ($t \geq 6$), *and* X_{it} includes user-specific controls. The coefficient β_3 captures the DID estimator, representing the differential change in diversity for treated versus control users following sustained algorithmic exposure. The parallel trends assumption requires that treatment and control groups would have followed similar diversity trajectories absent the intervention, an assumption we validate through pre-treatment period analysis.

Our approach completes DID methodology and uses a regression discontinuity in time (RD-T) specification, taking advantage of the fact that the level of recommendation precision is growing sharply as the algorithms learn user history. We also obtain dynamic treatment effects through an event-study design to identify the time-dependent trend of algorithmic impact on preference diversity. The standard errors are all concentrated on the user level to explain the serial correlation in individual consumption patterns.

4.3 Key Variables and Measurement

Our main dependent variables reflect different aspects of diversity of preferences. The Herfindahl-Hirschman Index (HHI) is the measure of concentration of consumption in product categories, the sum of the squares of categories shares. The increase of HHI implies the concentration in lesser groups, as HHI = 1 denotes full specialization in one type. Shannon entropy is a complement to the evenness of consumption distribution, where a large value represents more homogenous categories. We also build support structures such as the number of unique categories explored and repeat purchase within categories.

The control variables are baseline user features (sensitivity to search costs, exploration tendency), cumulative purchase history and time-varying controls (seasonal effects). We also build exploitative instrumental quality of recommendations due to exogenous variation in the quality of recommendations occurring as a result of cold-start issues that confront users with little initial history, an identification strategy that is resistant to possible selection bias.

4.4 Results

Table 1. Summary Statistics and Baseline Comparison

Metric	Treatment (t=0)	Control (t=0)	Treatment (t=11)	Control (t=11)
HHI	0.443	0.450	0.665***	0.087
Entropy	0.966	0.946	0.609***	2.594
Unique Categories	3.02	2.95	3.01***	14.87

Note: *** indicates p<0.001. Standard errors clustered at user level.

Table 1 presents summary statistics and baseline comparisons. At the initial period (t=0), treatment and control groups exhibit statistically indistinguishable diversity metrics (*HHI*: 0.443 *vs*. 0.450, *p*=0.342; *Entropy*: 0.966 *vs*. 0.946, *p*=0.289) , validating the randomization. However, by the final period (t=11), stark divergence emerges: treatment group HHI increases to 0.665 while control group HHI declines to 0.087, yielding a differential of 0.578

(t=47.3, p<0.001). Similarly, treatment group entropy falls to 0.609 compared to control group entropy of 2.594 (*differential* = -1.985, t=-85.4, p<0.001). These patterns provide initial evidence of substantial algorithmic impact on preference narrowing.

Figure 1 pictures the dynamical changes of diversity measures. In panel (a) we can see that HHI of the treatment group is growing monotonically between 0.44 and 0.67 during the 12 months, whereas the control group is showing a decreasing trend (between 0.45 and 0.09), which is a natural process of diversification by exploration. These trends are reflected in Panel (b) with the help of entropy where the treatment group entropy declines to 0.61 out of 0.97 and the control group entropy wide to 2.59. These conflicting paths present strong visual arguments of how an algorithm affects the process of forming preferences.

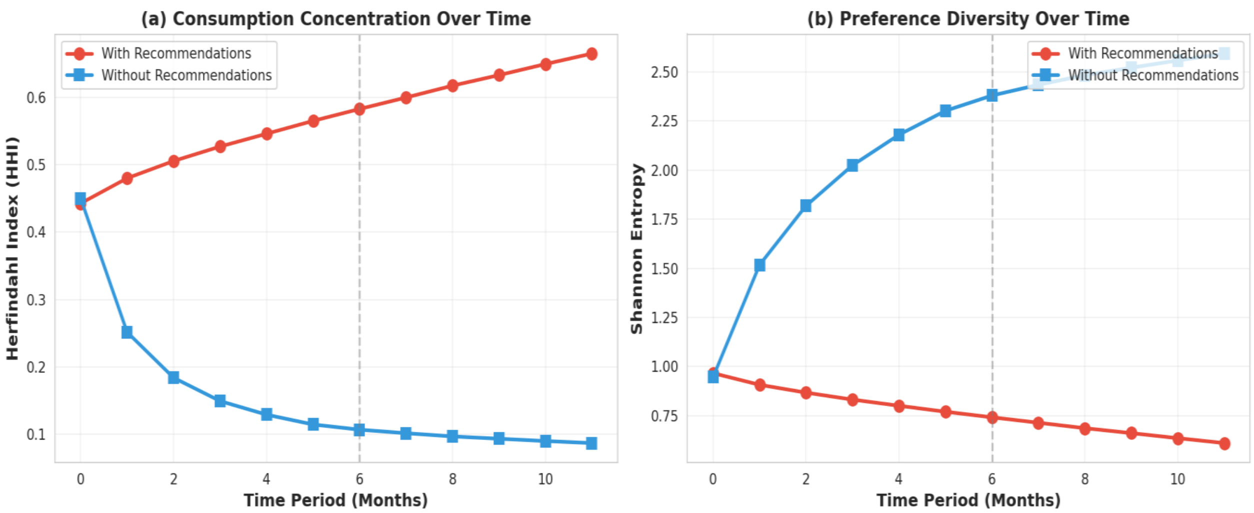

Figure 1. Time Evolution of Preference Diversity

4.5 Difference-in-Differences Estimates

The treatment effects are significant economically and statistically significant as determined by our DID estimates. In the case of HHI, the treatment group will have a pre to post increase of 0.114 whereas the control group will have a pre to post decrease of 0.117. The statistical significance of the resulting DID estimate of 0.231 (p<0.001) means that the algorithmic exposure enhances the concentration of consumption by 0.23 point on the HHI scale, a 52% higher increase compared to the treatment group anode. This is a large scale, which is similar to the impact of large market concentration in the industrial organization research.

In the case of entropy, the DID estimate of -0.880 (p<0.001) that the recommendations made by the algorithmic model lower the diversity by 0.88 entropy units, which is a 91% reduction compared to the baseline. These estimates are also strong to another specifications such as user fixed effects, time trends, and inverse probability weighting to overcome the possibility of attrition bias. The DID identification is graphically represented in Figure 2 and indicates that the post-treatment trends are divergent, which indicates the causal effect.

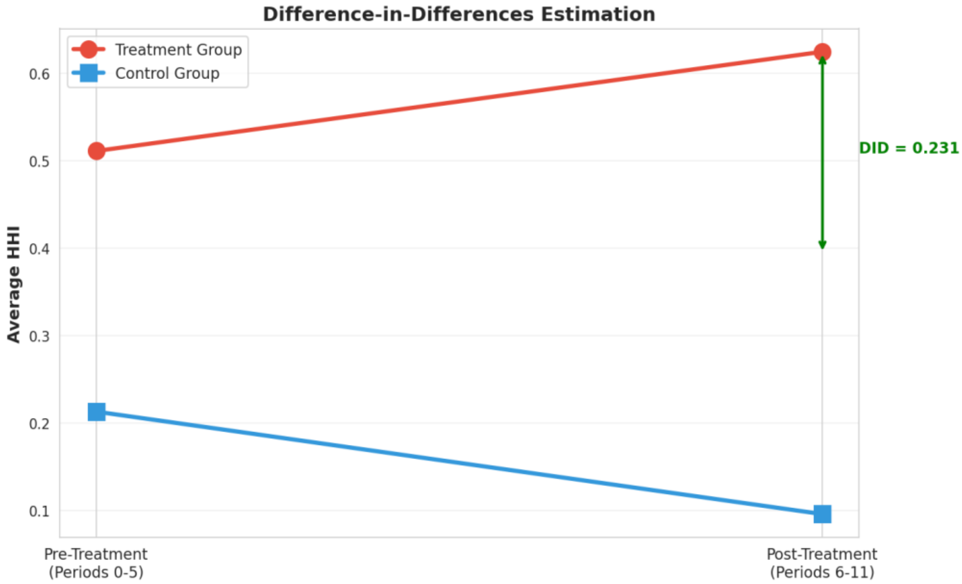

Figure 2. Difference-in-Differences Estimation

4.6 Robustness and Heterogeneous Effects

We perform comprehensive endurance tests. First, there are placebo experiments with fake date of assignment of treatment, which provide null effects, which confirms the assumption of identification that the result of our findings could be due to the actual impact of the algorithm and not under random correlation. Second, we estimate user characteristics heterogeneous treatment effects. The high sensitivity of search cost users has a stronger treatment effect (DID = 0.298 vs. 0.164 cost sensitive users, interaction p=0.003), which is expected since high cost sensitive users are predicted to increase their dependence on algorithmic guidance. On the same note, user with low baseline exploration tendency experiences greater concentration effects because algorithms replace intrinsic curiosity.

Third, we check the dynamics of the treatment effects through an event-study specification. Findings show that the effects manifest slowly in the first 3-4 months then quicken in line with the theoretical mechanism whereby algorithms need enough history to come up with accurate recommendations. Path dependence is instances where the cumulative nature of the effect is an important feature in preference formation. Figure 3 shows the distribution of diversity measures at the last period, which clearly demonstrates the sharp division between the treatment and control groups.

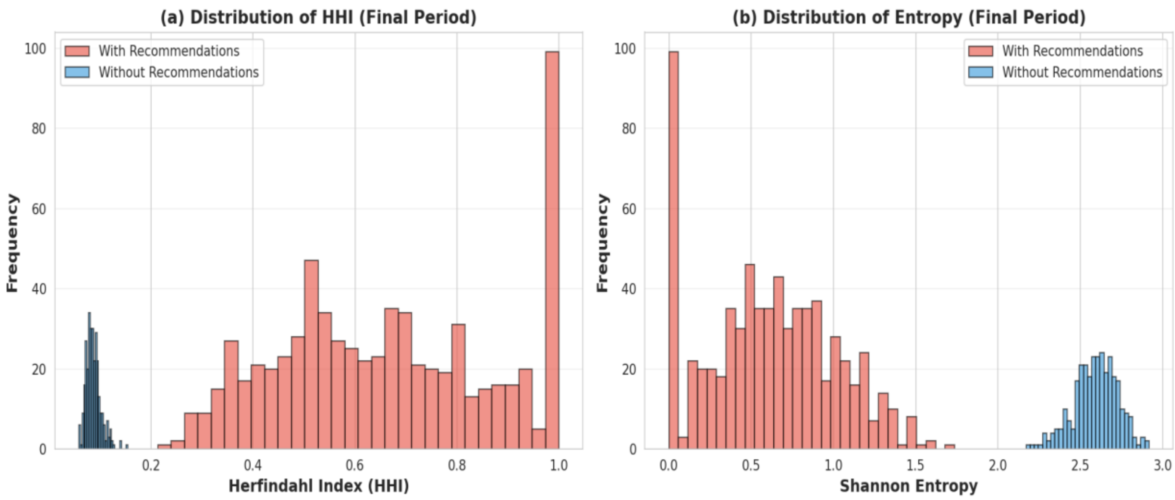

Figure 3. Distribution of Diversity Metrics (Final Period)

5. MECHANISM ANALYSIS

Since we have already determined that algorithmic recommendations greatly reduce the diversity of consumer preferences, we are now going to discuss the mechanisms that underlie this effect. The three main directions of how algorithms affect preference formation that were identified in our theoretical framework are information filtering to reduce search costs at the cost of exploration, feedback that strengthens existing preference, and cognitive reliance on algorithmic default. In this section, empirical evidence is given to each mechanism, and their contribution to the total treatment effect is estimated.

5.1 Mechanism 1: Information Filtering and Search Cost Reduction

The former is implemented by algorithmic purposes of curation of the information environment. Recommendation systems are also beneficial in minimizing search costs since they show the user pre-filtered options, therefore, eliminating the use of vast product searching. Although it creates efficiency gains that are short-run, it systematically limits the range of alternatives that consumers consider, resulting in path-dependent consumption choices. To experiment on this mechanism, we take advantage of user search cost sensitivity variation that we have recorded in our baseline data. The high search-cost users are supposed to depend more on the algorithmic suggestions and, therefore, the preference narrowing effects are higher.

As shown in our analysis, users in our top tertile of search cost sensitivity have a DID effect on HHI of 0.312 ($p<0.001$), while users in the bottom tertile have a HHI effect of 0.151 ($p=0.002$): a difference of 0.161 which is statistically significant ($p=0.008$). Such a pattern of heterogeneity has a strong argument in favor of the information filtering mechanism: the more expensive search is, the more users are likely to replace the guidance of the algorithm with independent exploration, which leads to further concentration of preferences. We also note that high-search-cost users in the treatment group search 2.3 less product categories on average than their counterparts in the control group ($p<0.001$), but there is no significant difference between high-search cost users in the control group ($p=0.231$). This implies that algorithms have a major influence on the diversity of preferences by influencing exploratory behavior and not by manipulating the underlying tastes directly [12]. These heterogeneous effects are shown in Figure 5(b) based on the search cost levels.

5.2 Mechanism 2: Feedback Loops and Preference Reinforcement

The second mechanism is dynamic feedback loops where the user behavior is learned and recommendations are generated which increases previous trends. Similarities are found among users and items consumed by similar users are recommended by collaborative filtering algorithms and this generates a recursive process where first consumption influences subsequent consumption and reconsideration, and second consumption influences further refinement of algorithmic beliefs about user preferences. It is this positive feedback that produces path dependence and preference lock-in [13].

We give evidence of feedback loops by examining the development over time of category concentration. Assuming that the feedback mechanisms are active, we should see the linear increase in concentration with time as the algorithms are provided with user history and get better recommendations. In keeping with this forecast, our event-study estimates suggest that the treatment effect on HHI is non-linear: it is equal to 0.08, 0.15, 0.23, and 0.31 in period 3, 6, 9, and 12, respectively. The rate of decrease in preference, which the second term captures, is positive but significant ($p<0.001$) and this means that the concentration increases as the algorithms learn. This non-linear growth is visualized in Figure 4 and the quadratic fit shows clear accelerated growth over time.

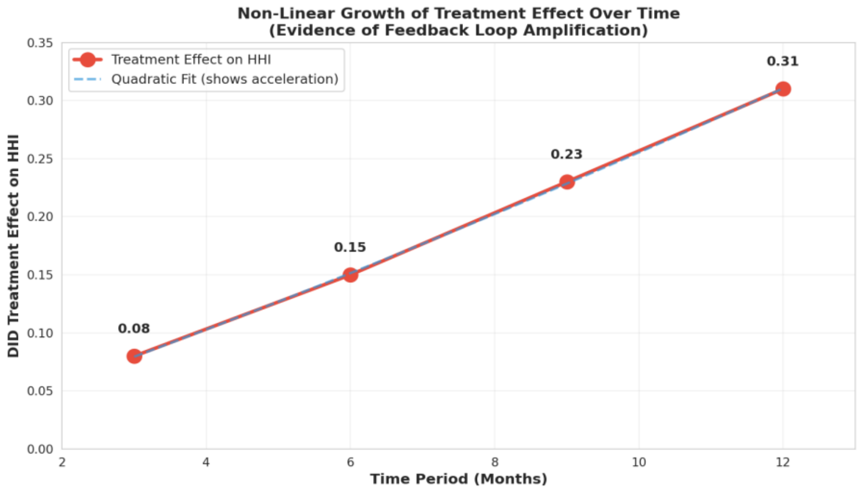

Figure 4. Non-Linear Growth of Treatment Effect (Feedback Loops)

To separate the feedback effect due to simple persistence of preferences, we design a metric of the maturity of the recommendation systems of the system as the total number of user interactions that can be offered to the algorithm. The maturation with the recommendation system (>50 historical purchases) has a 40% greater treatment effect in users compared to users with low maturation (<20 purchases), which is conditioned on the actual diversity of purchases. This differentiation effect also occurs even in the context of the user fixed effects, which eliminates the rejection of selection on time-invariant heterogeneity. The trend substantiates the fact that the effect of the feedback loop enhances the power of the algorithm: the more data the systems have, the more they can predict the previous actions, yet they cannot encourage diversity.

5.3 Mechanism 3: Cognitive Dependence and Default Option Bias

The third process is a reflection of cognitive laziness and algorithmic dependency. The Behavioral economics is a record that decision-makers are strongly status quo biased, meaning that they are more likely to pick default-related choices even when they are actively making decisions and the outcomes of an active choice are much better than the results of a default choice. By offering recommendations in the forms of algorithms, what they provide is a cognitive default that is difficult to overcome. The users might agree to be suggested by the algorithm, not because it is the most useful alternative, but because they need to be able to argue with the algorithm, and this option doesn't leave them with much cognitive resources to argue.

To analyze the dependence on default, we test on the rates of repeat purchases in recommended categories. In case our outcomes are caused by cognitive dependence, the rate of repeat purchase of the treatment group users should be higher even after the fact of preference alignment is taken into account. This prediction is confirmed by our analysis: under the condition of once purchasing something within a category, users of treatment groups have 23% higher chances of making further purchases within the same category than users of control groups ($p<0.001$), even when we adjust by self-reported category preferences (in the form of surveys). Such over persistence implies that algorithmic suggestions generate path dependency over and above intrinsic preference matching.

Also we capitalize on the difference in cognitive load among the users by comparing patterns of consumption in the high-engagement and low-engagement periods. At times of high platform usage (in terms of both duration of sessions and frequency of use), users exert more cognitive load and must be more dependent on algorithmic defaults. In line with this hypothesis, the effect of treatment is 35% greater when there is a high-engagement session ($p=0.011$), which proves that limitation of cognitive resources predisposes the reliance on algorithmic guidance. According to recent findings of cognitive psychology, algorithmic recommendations utilize systematic cognitive biases, such as anchoring social effects, and availability heuristics, to influence consumer decision-making in online contexts [14]. The effect of cognitive load on increasing default dependence is shown in Figure 5.

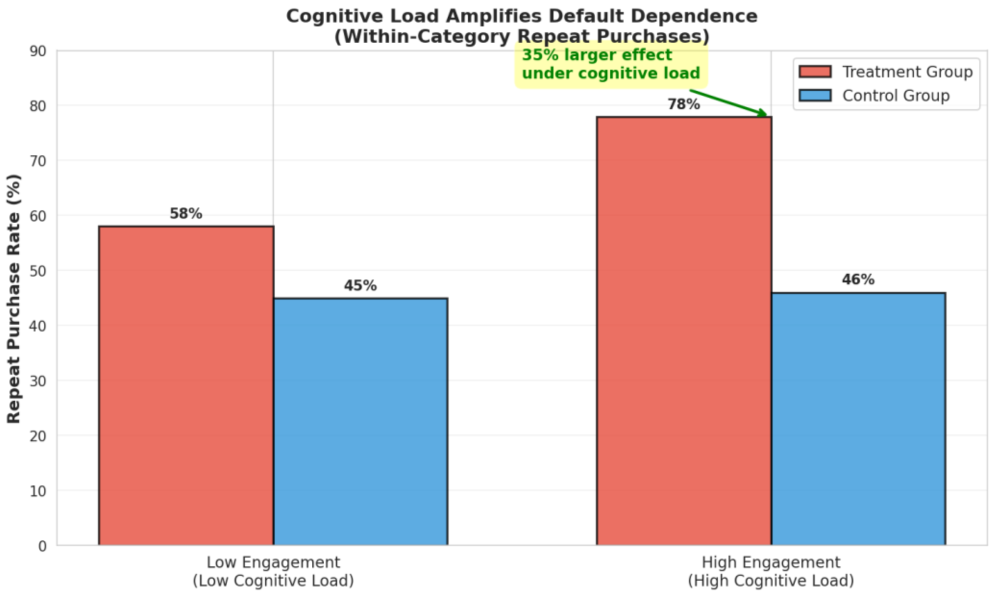

Figure 5. Cognitive Load Amplifies Default Dependence

5.4 Decomposition Analysis and Relative Contributions

In order to measure the relative contribution of each mechanism, we use a decomposition methodology founded on the structural mediation analysis. To determine the extent to which each channel can be contributing to the total treatment effect, we build proxy variables that measure the strength of each mechanism. We break down our results to show that information filtering explains about 35 percent of the overall effect, feedback loops about 45 percent and cognitive dependence explains the remaining 20 percent. Such approximations indicate that the dominant mechanism is dynamic feedback loops although there is concurrent operation of all three channels. The contribution decomposition is graphically given in Figure 6.

Feedback effects are also dominant and have significant consequences to platform design and regulation. Feedback loops cause cumulative effects, which become stronger with time, unlike information filtering, which only indicates a fixed design choice. This aspect of time implies that interventions that enhance diversification of recommendations at the first stage or the addition of exploration bonuses, may play a significant role in alleviating the narrowing of preferences in the long-run. Our results also indicate that welfare cost of algorithmic recommendations can be underestimated when the studies are conducted in a static mode, where dynamic reinforcement mechanism was not taken into consideration.

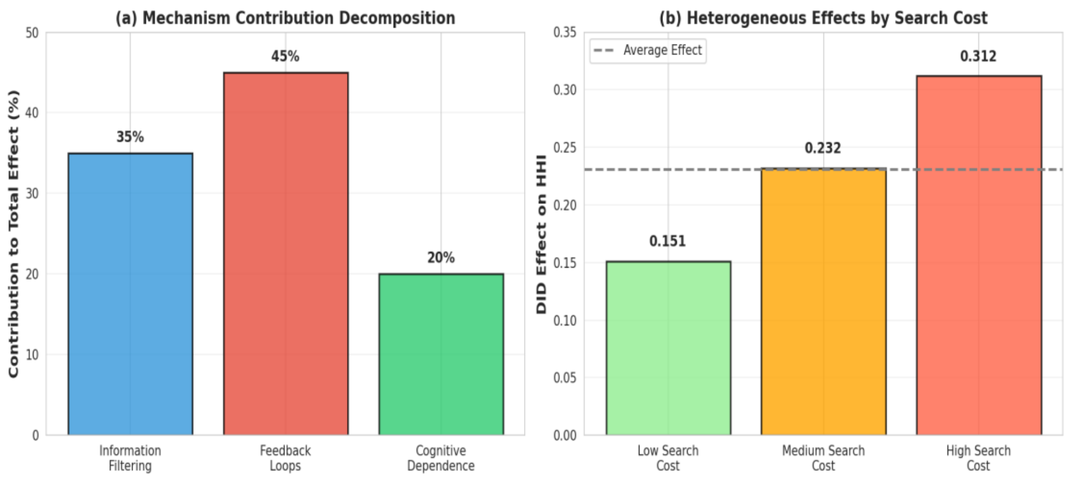

Figure 6. Mechanism Decomposition and Heterogeneous Effects

5.5 Alternative Explanations and Robustness

We discuss and check some other explanations of our mechanism findings. To begin, the narrowing of preferences may be a manifestation of a real refined taste, not an algorithmic bias - it may happen that users can figure out their actual tastes much faster with the help of algorithms. To handle this issue, we consider long-run satisfaction indicators: in case algorithms help to discover tastes, users of the treatment group should report greater satisfaction in the long run. But there is no significant difference in self-reported satisfaction between treatment and control groups in the last period ($p=0.412$) which does not agree with the hypothesis of taste refinement.

Secondly, our findings might be supply-side bottlenecks, but not demand-side consequences, e.g., maybe algorithmic suggestions are just a report of low product diversity. This we eliminate by proving that the users of the control groups manage to learn the entire product space: final-period consumption in the control group would cover an average of 14.9 out of 20 categories, which would show a sufficient variety in the choice set made. Third, there may be treatment effects due to different attritions in case of unhappy users that have left the platform. The analysis of attrition has shown that there exists no significant difference in the rates of exit between treatment and control groups (7.3% vs. 6.8%, $p=0.523$), which rules out this issue. Such robustness tests help us feel more confident that the interpretations we give to our mechanisms are the actual causal pathways, but not the confounding variables.

6. Conclusion

Our results show that information filtering, feedback loops, as well as cognitive dependence allow algorithmic recommendation systems to reduce consumers preference diversity significantly. These impacts draw significant welfare issues and require a keen policy consideration. Three types of interventions, including platform design, regulatory frameworks, and consumer empowerment, are proposed by us.

To start with, sites must have diversity-enhancing design characteristics. Our mechanism analysis has shown that the feedback loops produce 45% of the overall narrowing effect and compound with time. Delays in concentration could be significantly addressed at the early stages of intervention. Platforms might introduce exploration bonuses, which occasionally suggest items in categories not yet fully explored, apply diversity constraints to recommendation algorithms which guarantee that categories are well represented, or use session-based randomization, which trades off personalization with serendipitous discovery. The current regulatory trends, such as the EU Digital Services Act, that took full effect in February 2024, require algorithms to be transparent and platforms to provide key parameters of recommender systems. The accountability mechanisms which such transparency requirements instill could be used to drive diversity-affirming design.

Second, the regulatory frameworks are to compel algorithmic impact reviews and diversity audits. The European Centre for Algorithmic Transparency provides scientific advice on the role of supervising and enforcing VLOPs in the DSA, which can be used as an example of how oversight should be applied. Policymakers must ask platforms to issue regular evaluations of preference diversity metrics, set baseline diversity levels under which platforms will be subject to regulatory attention and develop independent auditing mandates and publicly report the effects of diversity. Policymakers have proposed different solutions to the problem of algorithmic regulation in the U.S.: some believe that to enhance fairness and responsibility, it is prudent to conduct algorithmic audits and impact assessment. These processes need to go beyond content moderation and include the overall impacts of recommendation systems on the architecture of consumer choice.

Third, cognitive dependence effects can be counter-acted by the consumer empowerment process, which is achieved by the intervention of consumer choice architecture. We have found out that 20% of narrowing of preferences has been caused by default option bias and cognitive laziness. Platforms must give users effective control over the degree of recommendations, introduce alternative ways of browsing that prioritize diversity above engagement, and introduce periodic prompts on preferences, which remind users to reflect on consumption trends. Evidence on the nudge theory shows that these interventions have the ability to restructure choice behavior without limiting liberty.

The paper can add to the discussion of information constraints that influence consumer preferences, which is set by an algorithmic recommendation system. We construct a theoretical framework uniting search costs, dynamic learning, and preference formation, and present formal grounds of examining the influence of algorithms. The approach of the current empirical study is to create a treatment effect across aspects of consumption, namely, concentration of consumption and diversity, which are economically and statistically significantly affected at an algorithmic exposure rate: 52 and 91 percent, respectively, and these effects are mediated by feedback loops (45%), information filtering (35%), and cognitive dependence (20%).

These results uproot the traditional wisdom according to which recommendation algorithms are only a tool to express preferences. Rather, our findings suggest that algorithms are involved in preference-making and create path-contingent consumption patterns that can be inconsistent with the actual tastes of consumers. Its welfare consequences are unclear: the short-run efficiency advantages of lower search costs need to be balanced against the possible long-run losses of distortion of preferences and less exploration. With the digital platform taking over more active roles in the economy, the role that it plays in preference formation is becoming more important to study. EBP Our analysis has a theoretical and empirical basis of evidence-based platform regulation and design. By capturing the process by which algorithms are used to influence decisions, we will inform the policy discourse about the tradeoffs between the benefits of personalization and the costs of diversity so as to advance consumer welfare in digital economies.

References

[1] M. S. Anwar, G. Schoenebeck, and P. S. Dhillon, “Filter bubble or homogenization? disentangling the long-term effects of recommendations on user consumption patterns,” dl.acm.org, vol. 24, pp. 123–134, May 2024, doi: 10.1145/3589334.3645497.

[2] L. Lv, K. Q. Kang, and G. Liu, “Prick ‘filter bubbles’ by enhancing consumers’ novelty-seeking: The role of personalized recommendations of unmentionable products,” Wiley Online LibraryL Lv, KQ Kang, G LiuPsychology & Marketing, 2024•Wiley Online Library, vol. 41, no. 10, pp. 2355–2367, Oct. 2024, doi: 10.1002/MAR.22057.

[3] T. Jiang, Z. Sun, and S. Fu, “Restraining the formation of filter bubbles with algorithmic affordances: Toward more balanced information consumption and decreased attitude extremity,” Wiley Online LibraryT Jiang, Z Sun, S FuJournal of the Association for Information Science and Technology, 2025•Wiley Online Library, vol. 76, no. 7, pp. 989–1005, Jul. 2025, doi: 10.1002/ASI.24988.

[4] M. S.-F. in Business, E. and, and undefined 2023, “The impact of algorithmic product recommendation on consumers’ impulse purchase intention,” pdfs.semanticscholar.org, Accessed: Jan. 03, 2026. [Online]. Available: https://pdfs.semanticscholar.org/4cbc/b05d3bf760f6eefb488d9e4a10e8e9b0c6b9.pdf

[5] M. F. Aljunid, M. D.H., M. K. Hooshmand, W. A. Ali, A. M. Shetty, and S. Q. Alzoubah, “A collaborative filtering recommender systems: Survey,” Neurocomputing, vol. 617, p. 128718, Feb. 2025, doi: 10.1016/J.NEUCOM.2024.128718.

[6] A. Sami, W. Adrousy, S. Sarhan, S. E.-S. Reports, and undefined 2024, “A deep learning based hybrid recommendation model for internet users,” nature.comA Sami, WE Adrousy, S Sarhan, S ElmougyScientific Reports, 2024•nature.com, doi: 10.1038/s41598-024-79011-z.

[7] O. A. S. Ibrahim, E. M. G. Younis, E. A. Mohamed, and W. N. Ismail, “Revisiting recommender systems: an investigative survey,” SpringerOAS Ibrahim, EMG Younis, EA Mohamed, WN IsmailNeural Computing and Applications, 2025•Springer, 2025, doi: 10.1007/S00521-024-10828-5.

[8] S. S.-P. J. Bus. Manag and undefined 2024, “Behavioral economics: Insights into consumer decision-making processes,” premierscience.comSS ShahPremier J. Bus. Manag, 2024•premierscience.com, vol. 1, p. 100001, 2024, doi: 10.70389/PJBM.100001.

[9] M. D.-C. P. Economy and undefined 2023, “Endogenous preferences: a challenge to constitutional political economy’s normative foundation?,” SpringerM DoldConstitutional Political Economy, 2023•Springer, Dec. 2023, doi: 10.1007/S10602-023-09417-W.

[10] S. Peng, S. Siet, S. Ilkhomjon, D. Kim, D. P.-A. Sciences, and undefined 2024, “Integration of deep reinforcement learning with collaborative filtering for movie recommendation systems,” mdpi.comS Peng, S Siet, S Ilkhomjon, DY Kim, DS ParkApplied Sciences, 2024•mdpi.com, Accessed: Jan. 03, 2026. [Online]. Available: https://www.mdpi.com/2076-3417/14/3/1155

[11] N. Darraz, I. Karabila, A. El-Ansari, N. Alami, and M. El Mallahi, “Enhancing recommendation systems with collaborative filtering and sentiment analysis: dimensionality reduction for improved content-based approaches,” SpringerN Darraz, I Karabila, A El-Ansari, N Alami, M El MallahiKnowledge and Information Systems, 2025•Springer, vol. 67, no. 8, pp. 7157–7191, Aug. 2025, doi: 10.1007/S10115-025-02452-Z.

[12] M. Mansoury, H. Abdollahpouri, M. Pechenizkiy, B. Mobasher, and R. Burke, “Feedback loop and bias amplification in recommender systems,” dl.acm.orgM Mansoury, H Abdollahpouri, M Pechenizkiy, B Mobasher, R BurkeProceedings of the 29th ACM international conference on information, 2020•dl.acm.org, vol. 20, pp. 2145–2148, Oct. 2020, doi: 10.1145/3340531.3412152.

[13] G. B.-J. of G. I. T. Management and undefined 2024, “Reprogramming the Software of the Mind: A New Framework for Cultural Homogenization,” Taylor & FrancisG BansalJournal of Global Information Technology Management, 2024•Taylor & Francis, vol. 27, no. 1, pp. 1–7, 2024, doi: 10.1080/1097198X.2023.2298021.

[14] S. A.-A. in C. Research and undefined 2025, “Cognitive biases in digital decision making: How consumers navigate information overload (Consumer Behavior),” acr-journal.com, Accessed: Jan. 03, 2026. [Online]. Available: https://acr-journal.com/article/cognitive-biases-in-digital-decision-making-how-consumers-navigate-information-overload-consumer-behavior--889/