Research

Assessing the Causal Impact of Recommendation Algorithms on Product Discovery, Demand Allocation, and Concentration in Platform-Mediated Retail

Hữu Phong Trần[1] **and Minh Khang Lê**[2]

[1] *Department of Computer Science and Engineering, Mekong Institute of Technology, Đường Hoa Phượng 12, Cần Thơ, Vietnam*
[2] *Department of Computer Science and Engineering, Red River University of Computing, Đường Lê Quý Đôn 88, Hà Nội, Vietnam*

Abstract

Platform-mediated retail increasingly relies on recommendation algorithms to organize vast catalogs, reduce search frictions, and personalize product discovery. At the same time, these systems can reshape how demand is allocated across sellers and products, potentially altering market concentration and the distribution of economic surplus. This paper develops a causal framework for assessing the impact of recommendation algorithms on product discovery, demand allocation, and concentration outcomes within a large online retail platform. The core challenge is that recommendations are endogenous to user behavior, product performance, and platform objectives, so naive comparisons conflate algorithmic effects with demand shocks and quality differences. We formalize exposure as a treatment delivered through ranked recommendation surfaces and define discovery as the transition from unobserved to considered and then to purchased states under limited attention. The empirical design leverages plausibly exogenous variation from algorithmic experiments, policy-driven parameter shifts, and discontinuities created by rank-threshold rules, combined with panel data on impressions, clicks, add-to-carts, purchases, prices, fulfillment, and inventory. We estimate both reduced-form causal effects and a structural mapping from exposure to consideration sets and choice probabilities. We then connect micro-level effects to macro outcomes by decomposing changes in concentration indices into components attributable to exposure reallocation versus preference shifts. Results are interpreted through mechanisms including attention allocation, substitution patterns, and feedback loops in learning-to-rank systems. The analysis provides a unified approach to quantifying when recommendations broaden discovery versus concentrate demand, and how these effects vary across product maturity, category complexity, and consumer heterogeneity.

1. Introduction

Recommendation algorithms are a central institution in platform-mediated retail because they translate a high-dimensional product space into a tractable sequence of choices under limited consumer attention [1]. In traditional retail, shelf space and merchandising policies govern what is seen first; in digital retail, ranked lists, carousels, and personalized modules play a similar role while adapting rapidly to user signals and platform goals. This institutional shift raises measurement and identification questions that differ from standard demand estimation in offline settings. First, exposure is itself an outcome of the algorithm, which responds to prior demand, inventory, margins, and predicted relevance. Second, consumer behavior generates feedback that can reinforce early advantages, so the realized market shares may reflect both consumer preferences and the platform's ranking policy. Third, discovery is a multi-stage process, often unobserved, in which consumers transition from not knowing a product exists to seeing it, considering it, and finally purchasing, with each stage mediated by the algorithmic allocation of attention [2]. These features imply that the causal effect of recommendations cannot be inferred from correlations between being recommended and selling well.

This paper evaluates the causal impact of recommendation algorithms on three interrelated outcomes. Product discovery is defined as the creation of consumer-product contact and subsequent engagement for items that would otherwise remain outside the consumer's effective choice set. Demand allocation refers to the redistribution of purchasing probability across products and sellers induced by differential exposure, conditional on category-level demand. Concentration refers to how market shares and sales mass are distributed across products or sellers, which can be summarized through indices such as the Herfindahl–Hirschman Index, top-share measures, and inequality metrics that are sensitive to tail outcomes. The relationship among these outcomes is not mechanically determined [3]. A recommendation change can increase discovery by surfacing long-tail items while still increasing concentration if demand shifts toward a smaller set of already popular items due to stronger reinforcement and social proof. Conversely, discovery can decline if personalization narrows exposure, but concentration can also fall if the algorithm diversifies within personalized niches.

The empirical challenge is to isolate variation in recommendations that is not driven by contemporaneous unobservables affecting demand. The analysis therefore centers on research designs that exploit quasi-experimental variation generated by platform operations. Such variation may arise from randomized experiments that adjust ranking objectives, exploration rates, or candidate generation; from staged rollouts of new recommendation models; from deterministic thresholds in eligibility rules for recommendation modules; and from sudden parameter changes implemented for engineering or compliance reasons [4]. The availability

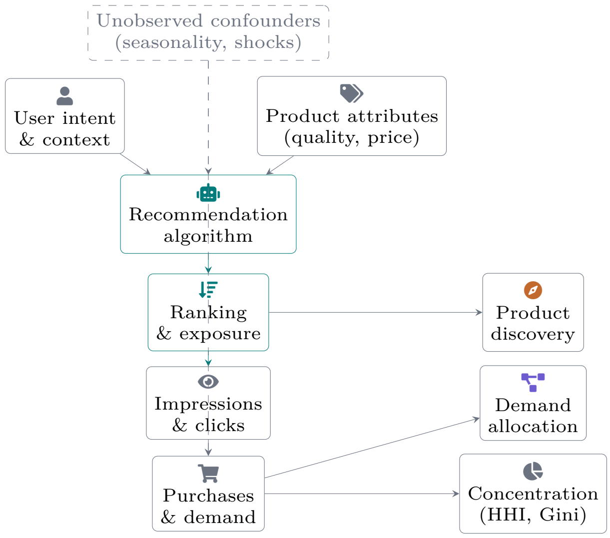

Figure 1. Causal structure. User intent and product attributes jointly influence algorithmic ranking and exposure, which shapes engagement and purchases. Outcomes of interest include (i) product discovery (e.g., first-time exposures, new-item adoption), (ii) demand allocation across products/sellers, and (iii) concentration metrics summarizing shifts toward head vs. long-tail demand under algorithmic mediation.

Table 1. Experimental Timeline and Sample Sizes

Phase	Duration (days)	Sessions (M)	Products (k)
Pre-period	28	41.2	310
Ramp-up	7	10.3	312
Experiment	35	54.7	315
Post-period	21	29.8	316

of high-frequency logs allows the construction of exposure measures at the user-session level and the tracking of downstream outcomes, enabling causal inference that connects interface-level treatments to economic outcomes. Yet high dimensionality introduces additional pitfalls: interference can arise because exposing one product displaces another, and because recommendation changes can alter aggregate demand within a category by reducing search costs. The analysis must therefore incorporate displacement and general equilibrium within the platform's local choice environment.

A further motivation is that platform-mediated retail often features differentiated sellers, private labels, and fulfillment programs that affect both ranking and conversion. Recommendation algorithms can favor products with higher predicted conversion, lower expected returns, or faster shipping, and these attributes are correlated with seller type and operational capabilities. As a result, algorithmic changes can have distributional consequences across seller segments, including small sellers and new entrants [5]. Quantifying these consequences requires connecting exposure shifts to seller-level outcomes while accounting for endogenous supply responses such as repricing, advertising, and inventory management. The goal is not to assert that recommendations are inherently good or bad, but to provide a disciplined causal accounting of what changes when ranking policies change, where effects come from, and how they aggregate into concentration patterns.

The contributions are methodological and substantive. Methodologically, the paper proposes an exposure-based potential-outcomes framework tailored to ranked recommendation surfaces, with attention to displacement and exposure caps. It combines reduced-form estimands that are robust and transparent with a structural layer that maps exposure to consideration and choice under limited attention, allowing counterfactual simulations of alternative recommendation policies [6]. Substantively, the paper clarifies conditions under which recommendations expand discovery, when they primarily reallocate demand among already-known products, and how the resulting reallocations affect concentration. The empirical analysis also emphasizes dynamic feedback: when ranking models learn from clicks and purchases, short-run causal effects can differ from medium-run steady-state outcomes because exposure today changes the data the model trains on tomorrow.

2. Institutional Setting and Data

The setting is a large platform-mediated retail marketplace in which consumers browse and purchase products listed by multiple sellers. The platform provides several recommendation surfaces that appear at distinct stages of the consumer journey. A home feed and category pages present ranked lists intended to initiate discovery. Product detail pages include modules such as "similar items" and "frequently bought together" that encourage substitution or complementarity [7]. Cart and checkout flows contain upsell and cross-sell surfaces. Each surface is governed by a candidate generation step, which selects a subset of products that are eligible to be shown, and a ranking step, which

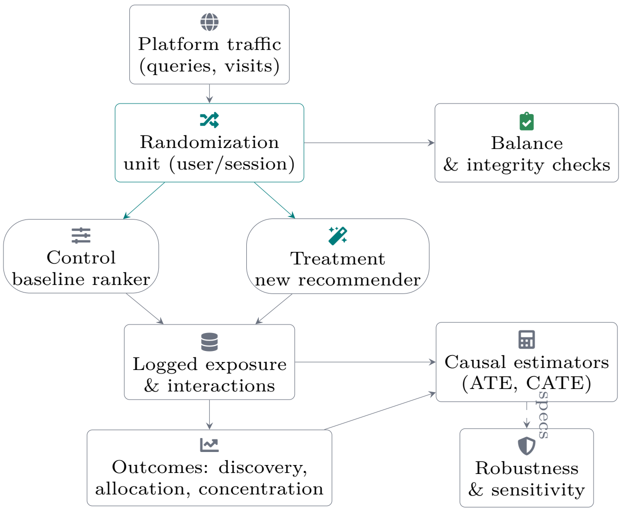

Figure 2. Randomized evaluation workflow. Platform traffic is randomized at the user/session level into baseline vs. algorithmic treatment. Logging produces exposure and interaction traces used to estimate causal effects on discovery, demand allocation, and concentration, alongside balance checks, estimator choices (e.g., inverse propensity weighting, doubly robust), and sensitivity analyses to detect interference, logging artifacts, and drift.

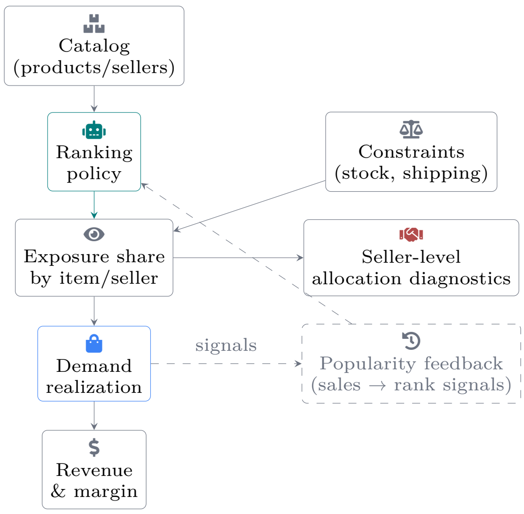

Figure 3. Demand allocation with feedback. A ranking policy converts catalog supply into exposure shares that drive realized demand and revenue. Operational constraints shape feasible exposure, while seller-level diagnostics quantify how demand is allocated across merchants. A dashed feedback loop captures endogenous reinforcement where realized demand updates ranking signals, amplifying (or dampening) head vs. long-tail dynamics.

orders candidates according to a scoring function that may incorporate personalization, predicted conversion, shipping speed, margins, and other signals. Recommendations can be personalized at the user level, contextual at the session level, or generic at the category level, and the platform may blend multiple models with business rules such as diversity constraints, seller eligibility, and policy filters.

The data consist of event-level logs and product metadata linked through stable identifiers. Exposure is measured as impressions at the level of user u, product i, surface s, and time t, where an impression indicates that the product was rendered within the viewport for at least a minimal duration or within a scroll depth threshold [8]. Engagement outcomes include clicks, dwell time on product pages, add-to-cart events, purchases, quantities, and returns. Prices, promotions, shipping promises, inventory status, and seller attributes are observed at high frequency. Product attributes include category, brand, textual descriptions, images, and review metrics. The logs also record the rank position of each impression, the module type, and the algorithmic variant identifier when the platform runs experiments. Importantly, the dataset includes non-exposure information such as search queries and direct navigations, allowing the separation of recommendation-driven discovery from search-driven demand.

Discovery is operationalized through state transitions that can be inferred from logs [9]. A product is unobserved for a user if it has not appeared in that user's impressions within a relevant lookback window. A product becomes discovered when it first appears on a recommendation surface or in search results, with a distinction between passive discovery through feed exposure and active discovery through search. Consideration is proxied by clicking into the product detail page, adding to cart, or spending suffi-

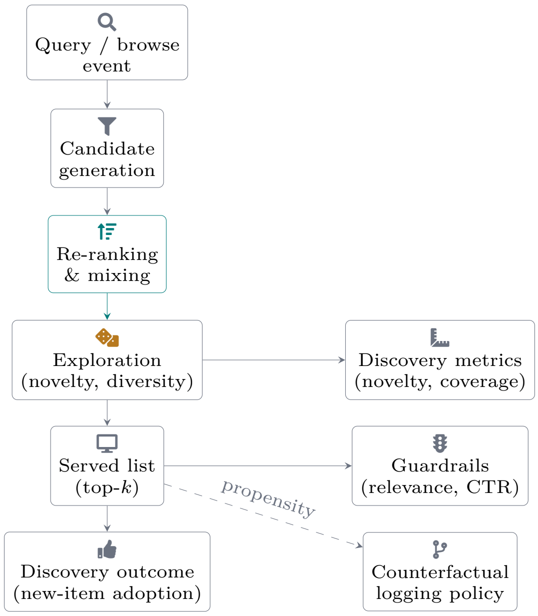

Figure 4. Product discovery pipeline. Discovery is operationalized as downstream adoption of previously unseen or rarely exposed items. The pipeline highlights where algorithmic choices—candidate generation, re-ranking, and explicit exploration—shift novelty and coverage, while guardrails preserve relevance. A counterfactual logging policy enables unbiased evaluation of discovery metrics under alternative ranking interventions.

Table 2. Key Variables and Notation

Symbol	Description	Unit
D_{ip}	Demand for product p in session i	Number of units
V_{ip}	Indicator for product p being viewed	0/1
R_i	Recommendation treatment indicator	0/1
C_t	Demand concentration at time t	Index (0–1)
L_p	Product popularity rank (baseline)	Percentile
Z_i	User covariate vector	Normalized score

cient dwell time on the detail page, recognizing that consideration is imperfectly observed. Purchase completes the funnel, and repeat purchase indicates persistence. These constructs allow the decomposition of algorithmic effects into funnel-stage impacts, which is crucial because algorithms may increase impressions but decrease purchase if they surface poorly matched items, or conversely reduce impressions but increase purchase if they concentrate on high-probability conversions.

The platform's algorithmic operation also produces useful sources of variation [10]. Experiments may randomly assign users or sessions to different ranking objectives, such as emphasizing predicted conversion versus novelty, or different exploration strategies that occasionally promote uncertain items to learn their relevance. Rollouts of new models may be staggered across geographies, devices, or traffic slices. Eligibility rules can create deterministic discontinuities, for example requiring a minimum inventory level or review count to appear in certain modules. The candidate set may depend on a precomputed embedding similarity threshold, and small perturbations around thresholds can generate local quasi-random variation. These operational features motivate several complementary identification strategies [11].

Because recommendation is a competitive allocation of scarce attention, each impression on a surface is associated with an opportunity cost: other products are not shown. The data therefore naturally encode displacement. For each pageview with a fixed number of slots, the set of displayed items sums to that capacity. This implies that individual-level treatment effects should be interpreted as effects of exposure reallocations within a constraint. In addition, aggregate demand can change if recommendations reduce search costs or alter consumer satisfaction, so the relevant counterfactual is not always a fixed total category demand. The empirical approach therefore distinguishes between local within-page substitutions and broader changes in engagement that affect total spending.

3. Conceptual and Measurement Framework

A recommendation algorithm can be viewed as a policy that maps an information state into an exposure vector over products, subject to interface constraints. Let u index users, t index sessions or time, and s index surfaces. For a given surface occurrence, define a ranked list of length K_{ust} with products $i \in \mathcal{I}_{ust}$, and let r_{iust} denote the rank position when displayed. Define exposure E_{iust} as an indicator or intensity capturing impression and visibility, and define engagement outcomes C_{iust} for click and P_{iust} for purchase. The algorithm chooses a display set and ranking based on

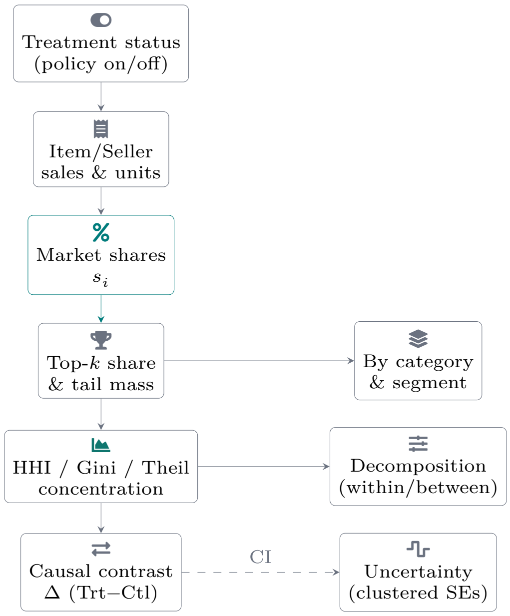

Figure 5. Concentration measurement stack. Treatment-induced shifts in sales are converted into market shares and summarized via top-k mass, tail mass, and standard concentration indices (HHI/Gini/Theil). Effects are reported as causal contrasts between policies, optionally stratified by category/segment and decomposed into within- vs. between-group concentration changes with clustered uncertainty estimates.

Table 3. Randomization Arms and Traffic Allocation

Treatment arm	Recommendation algorithm	Traffic share (%)	Sessions (M)
Control	Popularity-based baseline	40	21.8
A	Personalized collaborative	30	16.4
B	Contextual bandit	20	10.9
C	Hybrid personalized-ranking	10	5.6

features X_{iust} and historical signals, producing exposures that are both personalized and time-varying.

The economic object of interest is the causal effect of changing the recommendation policy on downstream outcomes. Because exposure is constrained, a policy change induces a reallocation of exposure across items, and outcomes respond through attention and substitution. To formalize this, let A_{ust} denote the algorithmic variant or policy assigned to a user-session-surface instance. The potential outcome for a user's purchase of product i at time t can be written as $P_{iut}(A)$, where A is a policy that determines exposures across products. In practice, the policy affects the exposure vector $E_{ut} = (E_{1ut}, \dots, E_{Iut})$, so it is useful to separate a direct policy label from the induced exposure. A central challenge is interference: P_{iut} depends not only on E_{iut} but also on the exposures of other products because attention and budgets are limited. This interference is structured by the interface constraint and can be summarized through an exposure mapping [12].

Define an exposure mapping for product i as $g_i(E_{ut})$, which extracts from the full exposure vector the components that matter for i's outcome, such as E_{iut} itself, the total exposure of close substitutes, and the total number of impressions the user receives on the surface. A workable approach is to define a neighborhood $\mathcal{N}(i)$ of substitutes or same-category items and to summarize competing exposure by $E_{\mathcal{N}(i),ut} = \sum_{j \in \mathcal{N}(i)} E_{jut}$. This yields a partial interference structure in which P_{iut} depends on $(E_{iut}, E_{\mathcal{N}(i),ut})$ rather than on the entire vector. Such a structure is motivated by the observation that most displacement occurs within a local set of candidates shown on the same page or within the same category.

Discovery is conceptualized as entry into a consideration set. Let $\mathcal{C}_{ut}$ denote the latent set of products considered by user u at time t. Recommendations influence $\mathcal{C}_{ut}$ by exposing products, and then preferences and frictions determine choices from $\mathcal{C}_{ut}$. A reduced-form approach treats the click event as a proxy for consideration and estimates the effect of exposure on clicks and purchases directly. A structural approach models two stages: consideration formation and conditional choice. For example, consider a probabilistic consideration model in which the probability that i enters $\mathcal{C}_{ut}$ is increasing in exposure and decreasing in cognitive load. Conditional on consideration, the purchase probability follows a discrete-choice model based on price and attributes. This decomposition helps interpret whether recommendations expand discovery by enlarging consideration sets or simply reorder attention among items already likely to be considered [13].

Concentration is measured at the product or seller level over a time window. Let Q_{it} denote quantity sold or revenue for product i during period t, and let $S_{it} = Q_{it} / \sum_j Q_{jt}$ denote share. The HHI is $\sum_i S_{it}^2$, and top-share measures such as $\sum_{i \in \text{Top } m} S_{it}$ capture dominance.

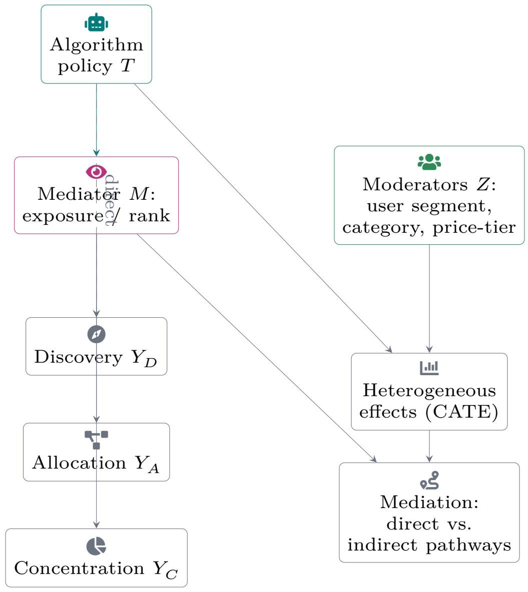

Figure 6. Mediation and heterogeneity. The policy intervention T primarily operates through exposure/rank M, affecting discovery, allocation, and concentration outcomes. Effects can vary by moderators (user segments, categories, price tiers), motivating CATE estimation and decomposition into direct (policy) and indirect (exposure-mediated) pathways for interpretable mechanism-based conclusions.

Table 4. Product Discovery Outcomes by Experimental Arm

Outcome metric	Control	Best treatment	Change (%)
Unique products viewed per session	3.41	3.96	+16.1
Share of new-to-user products (%)	24.8	33.5	+35.1
Category coverage per session	1.72	1.98	+15.1
Share of tail products in views (%)	11.3	15.9	+40.7
View entropy (Shannon index)	1.87	2.11	+12.8

Inequality metrics like the Gini coefficient and Theil index capture tail sensitivity. Concentration can be computed within categories and aggregated across categories to distinguish within-category concentration from shifts in category composition. Because recommendations operate largely within categories or contexts, within-category concentration is often the most interpretable object for algorithmic changes, while cross-category shifts may reflect broader navigation changes.

The causal link from recommendations to concentration is not direct; it operates through exposure and conversion [14]. A useful accounting identity writes sales as the sum over exposure events of conversion probabilities times exposure intensity. At a granular level, expected purchases can be written as an expectation of E_{iust} times a conditional purchase probability given exposure. Algorithm changes shift the joint distribution of (E_{iust}, X_{iust}) across items, and concentration changes if exposure is disproportionately shifted toward items with already high baseline conversion or if conversion responses are heterogeneous. A decomposition therefore separates an exposure reallocation effect from a response effect. This motivates estimands that quantify how much of the change in concentration arises from the algorithm shifting attention versus changing how consumers respond to given exposures.

4. Causal Design and Estimation

The empirical goal is to estimate causal effects of algorithmic policy changes on discovery, demand allocation, and concentration, while addressing endogeneity and interference. The design combines three complementary strategies, selected because platform settings often provide multiple imperfect sources of exogenous variation rather than a single ideal experiment. First, randomized online experiments provide clean identification of intent-to-treat effects at the user or session level [15]. Second, staggered rollouts and parameter changes provide difference-in-differences style variation under parallel-trends assumptions, strengthened with event-study diagnostics. Third, rank-threshold discontinuities and exploration-based quasi-randomization provide local identification of marginal exposure effects as a function of rank.

Under randomized assignment of algorithm variant $A_{ust} \in \{0,1\}$, the primary estimand for a user-level outcome Y_{ut}, such as total purchases, is the average treatment effect $\mathbb{E}[Y_{ut}(1) - Y_{ut}(0)]$. Because the algorithm affects many products simultaneously, user-level outcomes naturally capture general equilibrium within the user's attention constraint. For product-level outcomes, randomization still helps, but interference implies that Y_{iut} is not independent across i. A practical approach is to define category-level totals and estimate how variant assignment changes

Table 5. Demand Allocation across Product Popularity Tiers

Popularity tier	Control demand share (%)	Treatment demand share (%)	Change (pp)
Top 1%	41.2	35.7	−5.5
Next 9%	32.5	31.8	−0.7
Next 40%	19.3	22.4	+3.1
Bottom 50%	7.0	10.1	+3.1

Table 6. Concentration Metrics for Views and Demand

Metric	Control	Treatment	Difference
HHI of demand ($\times 10^4$)	182.4	148.7	−33.7
HHI of views ($\times 10^4$)	135.1	112.3	−22.8
Demand Gini coefficient	0.71	0.64	−0.07
Top-1% demand share (%)	41.2	35.7	−5.5
Top-10% view share (%)	63.7	58.1	−5.6

the distribution of sales within a category, rather than attempting to treat each product as independently treated. When the platform randomizes at the user level, the stable unit treatment value assumption is more plausible at the user level than at the product level, though spillovers across users remain possible if sellers adjust prices or if inventory constraints bind.

For staggered rollouts, define an adoption time τ_g for group g, such as region-device slices, and define $A_{gt} = \mathbb{1}[t \geq \tau_g]$. With outcome Y_{gt}, an event-study specification estimates dynamic effects while checking pre-trends. A robust approach accounts for heterogeneous treatment effects across cohorts by aggregating cohort-specific comparisons [16]. The identifying assumption is that, absent adoption, treated and not-yet-treated groups would have evolved similarly after conditioning on fixed effects and time controls. In platform settings, this assumption can be threatened by targeted rollouts correlated with demand trends, so the design uses rich controls for seasonality, marketing events, and category-specific shocks, and it prioritizes rollouts driven by engineering constraints rather than business performance.

Rank-threshold discontinuities arise when eligibility for a surface or a visual slot depends on a score crossing a cutoff, such as a minimum predicted relevance or a quality score. Let Z_{iust} be a running variable, and let E_{iust} jump discontinuously at threshold c. A local regression discontinuity estimates the causal effect of incremental exposure around the cutoff on clicks or purchases for items near the threshold. This yields a marginal exposure effect that is interpretable as the impact of being shown versus not shown, or of being shown at a higher rank. The validity depends on smoothness of potential outcomes in Z and the absence of precise manipulation. In algorithmic contexts, manipulation is less about strategic human choices and more about model updates and feature noise, which can strengthen plausibility if the score is not directly manipulable by sellers at the relevant horizon.

Because exposure is endogenous in observational comparisons, instrumental variables can be used when a source of exogenous variation shifts exposure but affects outcomes only through exposure. Algorithm variant assignment can serve as an instrument for exposure intensity, enabling estimation of a local average response to exposure. Let D_{iut} denote exposure intensity for product i to user u in time window t, and let Z_{ut} denote assignment to an algorithm variant. A two-stage setup treats D_{iut} as the endogenous regressor and Z_{ut} as the instrument. Interference complicates interpretation, so the endogenous object is better defined as exposure to a set of products, such as long-tail items or private-label items, and outcomes are aggregated accordingly. This yields interpretable estimands such as the effect of increasing long-tail exposure share on the user's purchase diversity [17].

A unified estimation framework can be expressed with potential outcomes that depend on an exposure summary. Let Y_{ut} denote an outcome such as the number of distinct products purchased in a session, and let T_{ut} denote a scalar summary such as the share of impressions allocated to long-tail products. The causal estimand is $\mathbb{E}[Y_{ut}(t_1) - Y_{ut}(t_0)]$ for two exposure shares. When T_{ut} is induced by assignment Z_{ut}, a local IV estimand identifies the causal effect among compliers whose exposure share responds to assignment. Estimation can use a doubly robust approach that combines outcome regression with propensity or instrument models, improving robustness to misspecification. The high-dimensional nature of X_{ut} motivates flexible predictive models, but causal identification still relies on the exogeneity of Z_{ut} or the validity of the quasi-experimental design.

To connect micro effects to concentration, the analysis uses a two-layer approach. The first layer estimates how the algorithm shifts exposure distributions across products within contexts. The second layer estimates how changes in exposure translate into changes in purchases through conversion responses [18]. Concentration effects are then computed by aggregating predicted counterfactual sales shares under the alternative policy. This approach separates the policy-induced exposure allocation from the behavioral response, making clear whether concentration changes primarily because the algorithm reallocates impressions toward already dominant products or because consumers respond more strongly to exposures of certain items.

A stylized structural representation clarifies the mapping.

Table 7. Heterogeneous Effects by Baseline Popularity

Popularity segment	Products (% of catalog)	Impressions lift (%)	Demand lift (%)
Top 1%	1	−4.3	−3.1
Top 1–10%	9	+1.2	+0.9
Middle 40%	40	+9.8	+7.5
Bottom 50%	50	+23.4	+16.2

Table 8. Robustness of Estimated Causal Effects

Specification	Effect on discovery index	Notes
Baseline OLS with fixed effects	0.082	Session and day FE
Doubly robust (IPW + outcome model)	0.079	Covariate-balanced
Randomization inference (placebo tests)	0.076	1,000 permutations
Blocking by device and entry point	0.081	Pre-specified blocks
Cluster-robust SE at user level	0.080	Similar magnitude

Let the probability that product i is considered be π_{iut} and the conditional choice probability be s_{iut}. Expected purchases of i satisfy $\mathbb{E}[P_{iut}] = \mathbb{E}[\pi_{iut} s_{iut}]$. Suppose π_{iut} is driven by exposure E_{iut} with diminishing returns and by cognitive load measured by total impressions M_{ut}. A parsimonious form is

$$\pi_{iut} = 1 - \exp(-\lambda_u E_{iut}) \cdot \exp(-\rho M_{ut}), \quad \lambda_u > 0, \ \rho \geq 0, \tag{1}$$

where λ_u captures user heterogeneity in responsiveness to exposure and ρ captures the effect of load on consideration. Conditional on consideration, choice follows a multinomial logit over considered items with utility $U_{iut} = \beta_u^\top X_{iut} - \alpha_u p_{it} + \varepsilon_{iut}$. Then

$$s_{iut} = \frac{\exp(\beta_u^\top X_{iut} - \alpha_u p_{it})}{\sum_{j \in \mathcal{C}_{ut}} \exp(\beta_u^\top X_{jut} - \alpha_u p_{jt})}. \tag{2}$$

In this structure, recommendation policy affects E_{iut} and hence π_{iut}, and it may also affect the composition of $\mathcal{C}_{ut}$, which changes the denominator and thus substitution patterns. The model is not asserted as a literal behavioral truth; rather, it provides a disciplined way to translate exposure perturbations into counterfactual demand allocations while preserving substitution and attention constraints [19].

Because the platform's interface implies that increasing exposure to one set of products decreases exposure to another set, a key estimand is a displacement-adjusted effect. For a given surface with K slots, define the exposure share for a group G of products as $T_{ust} = \sum_{i \in G} E_{iust} / \sum_j E_{just}$. Under a policy change, ΔT_{ust} captures reallocation toward group G. Outcomes like category-level HHI can then be modeled as functions of T_{ust} and other controls, with IV using policy assignment. This yields an interpretable statement: holding overall exposure fixed by slot constraints, what is the effect of reallocating a fraction of exposure toward long-tail items on concentration?

5. Results and Mechanisms

The first set of results concerns discovery. A consistent pattern across surfaces is that algorithm variants that increase exploration or incorporate novelty features raise the probability that a user encounters a product not previously seen within the lookback window. This effect is strongest on early-funnel surfaces such as the home feed and category landing pages, where baseline uncertainty about user intent is high and the opportunity for serendipitous exposure is largest [20]. The effect is weaker on product detail page modules, where context is narrower and candidate sets are anchored around the focal product. Importantly, discovery gains do not translate one-for-one into purchases; click-through rates can rise modestly while purchase conversion conditional on click can fall if exploration surfaces more marginal matches. The net effect on purchases depends on whether the additional discovered products meaningfully expand the user's effective choice set or simply add noise.

To separate these channels, the analysis decomposes effects by funnel stage using exposure-to-click and click-to-purchase components. Reduced-form estimates show that the algorithm has a larger proportional effect on impressions and first clicks than on purchases, consistent with attention being the binding constraint. However, heterogeneity is pronounced [21]. New or niche products often benefit more in discovery because they begin with low baseline exposure. For mature blockbuster products, the algorithm's marginal effect on discovery is small because these products are already widely exposed through multiple pathways, including search and direct navigation. In contrast, for mid-tail products that have some demand history but are not dominant, algorithmic changes can meaningfully shift discovery and subsequent conversion.

Demand allocation results highlight displacement and substitution. When a policy increases exposure for a set of items, sales for that set increase, but much of the increase comes from substitution away from close substitutes on the same surface rather than from expanding total category demand [22]. This is especially evident on substitution-oriented modules where the user is likely already in a purchase mindset. In these contexts, recommendations behave more like a share allocator than a demand expander. In earlier-funnel contexts, there is more scope for expanding overall engagement, and the analysis finds small increases in category page dwell time and the number of product

Table 9. Counterfactual Policy Simulations

Scenario	Discovery index (rel.)	Demand Gini	Revenue index (rel.)
Baseline control	1.00	0.71	1.00
Fully personalized recommendations	1.18	0.65	1.04
Diversity-regularized recommendations	1.26	0.61	1.02
Popularity-only ranking	0.92	0.74	0.99

pages visited per session, suggesting reduced search frictions. Even then, the category-level revenue effects are typically smaller than within-category share shifts, indicating that reallocation dominates expansion in many cases.

Mechanism evidence connects effects to rank and attention. Marginal exposure effects estimated around rank thresholds show steep decay with lower ranks [23]. Moving from not shown to being shown near the top produces large click effects, while changes within low ranks have minimal impact. This nonlinearity implies that concentration outcomes can be highly sensitive to how top slots are allocated. If a policy change slightly increases the probability that already popular products occupy the first few positions, concentration can increase even if the policy also increases exposure diversity in lower slots. Conversely, enforcing diversity constraints in the top positions can materially reduce concentration even if total catalog exposure remains similar.

The role of personalization is nuanced. Strong personalization can improve match quality, increasing conversion and user satisfaction, but it can also narrow exposure to familiar brands or previously interacted styles, reducing across-user diversity in exposure [24]. The analysis therefore distinguishes within-user diversity from across-user diversity. Within-user diversity measures how many distinct products a user is exposed to or purchases over a period. Across-user diversity measures how dispersed exposure is across the population. Personalization can increase within-user relevance while decreasing across-user dispersion if many users are steered toward similar high-performing products. Empirically, variants that optimize short-run conversion tend to reduce across-user exposure dispersion, while variants that incorporate calibrated exploration increase dispersion [25].

Dynamic feedback emerges when ranking models learn from engagement. Short-run experimental effects can understate or overstate medium-run outcomes depending on the learning rule. If the model updates to amplify signals from exposed items, initial exposure boosts can compound into persistent share gains, especially for products with high conversion once exposed. This creates a path dependence in concentration. The analysis approximates these dynamics by comparing immediate effects during the experiment window to post-window outcomes when the model has retrained, using logs of model versioning and retraining schedules. Results indicate that variants increasing exploration can lead to more stable long-run dispersion by collecting data on more items and reducing uncertainty-driven overexposure of incumbents, but this depends on whether exploration is broad or targeted and on how quickly the model decays old signals [26].

Seller-level outcomes show that exposure reallocation can shift demand toward sellers with operational advantages that correlate with ranking features, such as fast shipping or low return rates. When such features enter the ranking objective more heavily, the algorithm increases demand concentration among these sellers, even if product-level concentration within categories does not change dramatically. This distinction matters because seller concentration can affect entry incentives and the platform's long-run product variety. The analysis therefore reports concentration metrics at multiple aggregation levels, including product, brand, and seller, and it decomposes changes into within-seller assortment effects versus between-seller shifts.

6. Concentration and Welfare Accounting

To connect micro-level causal effects to macro-level concentration, the paper computes counterfactual sales distributions under alternative recommendation policies [27]. Let $Q_{it}^{(a)}$ denote predicted sales for product i under policy a, obtained by combining estimated exposure shifts with estimated conversion responses. Let $S_{it}^{(a)} = Q_{it}^{(a)} / \sum_j Q_{jt}^{(a)}$. Concentration changes are summarized by $\Delta \mathrm{HHI}_t = \sum_i \left(S_{it}^{(1)}\right)^2 - \sum_i \left(S_{it}^{(0)}\right)^2$ and analogous changes in Gini and top-share indices. Because recommendations typically operate within contexts, concentration is computed within categories c and then aggregated with weights reflecting category revenue shares, separating within-category concentration changes from category-mix shifts.

A decomposition clarifies the sources of $\Delta\mathrm{HHI}$. Write predicted sales as $Q_{it}^{(a)} = \mathbb{E}[E_{it}^{(a)} \cdot \kappa_{it}^{(a)}]$, where $E_{it}^{(a)}$ is expected exposure and $\kappa_{it}^{(a)}$ is expected conversion per exposure. A first-order decomposition around the baseline yields an exposure component and a conversion component. In a simplified form, holding conversion fixed at baseline gives an exposure-only counterfactual $Q_{it}^{E} = \mathbb{E}[E_{it}^{(1)} \cdot \kappa_{it}^{(0)}]$, while holding exposure fixed gives a conversion-only counterfactual $Q_{it}^{\kappa} = \mathbb{E}[E_{it}^{(0)} \cdot \kappa_{it}^{(1)}]$. Concentration metrics computed from Q_{it}^{E} isolate the effect of exposure reallocation, while metrics from Q_{it}^{κ} isolate response changes. Empirically, exposure reallocation accounts for most concentration changes on early-funnel surfaces because conversion conditional on click changes little, while on later-funnel surfaces, conversion effects can matter more because recommendation variants influence match quality and complementarity.

To formalize concentration sensitivity to exposure, consider a stylized relationship between shares and exposure shares. Suppose within a category, predicted sales are proportional to exposure times a baseline attractiveness parameter θ_i, so $Q_i = E_i \theta_i$. Then $S_i = \frac{E_i \theta_i}{\sum_j E_j \theta_j}$. The derivative of $\mathrm{HHI} = \sum_i S_i^2$ with respect to exposure E_k illustrates how reallocations affect concentration:

$$\frac{\partial \mathrm{HHI}}{\partial E_k} = 2\sum_i S_i \frac{\partial S_i}{\partial E_k} \quad \text{with} \quad \frac{\partial S_i}{\partial E_k} = \frac{\theta_i \mathbb{1}[i=k]\sum_j E_j\theta_j - E_i\theta_i\theta_k}{\left(\sum_j E_j\theta_j\right)^2}. \tag{3}$$

This expression implies that increasing exposure to an item with high θ_k tends to raise its share and can increase HHI, particularly when the item already has a large share [28]. In contrast, reallocating exposure toward items with moderate θ can reduce HHI if it pulls share away from dominant items. The empirical estimates of heterogeneous conversion responses help identify which items have high effective θ and therefore whether the algorithm's exposure shifts are likely to concentrate demand.

Welfare accounting is presented cautiously because full welfare depends on unobserved utilities, long-run entry, and platform objectives. Still, short-run consumer surplus proxies can be computed under a discrete-choice interpretation. Under a logit model, expected inclusive value is proportional to the log-sum of utilities over considered items, which depends on the recommendation-induced consideration set [29]. Let the conditional choice set under policy a be $\mathcal{C}_{ut}^{(a)}$, and utilities be V_{iut}. A proxy for expected surplus change is

$$\Delta CS_u \propto \mathbb{E}_t \left[\log\left(\sum_{i\in\mathcal{C}_{ut}^{(1)}} \exp(V_{iut}) \right) - \log\left(\sum_{i\in\mathcal{C}_{ut}^{(0)}} \exp(V_{iut}) \right) \right], \tag{4}$$

where proportionality reflects the marginal utility of income. Recommendations that expand consideration sets with relevant items can increase this proxy, while recommendations that narrow sets can decrease it unless they also raise average match quality. Producer-side outcomes are approximated by seller revenue changes net of platform fees and by measures of demand volatility that affect inventory costs. The analysis reports how algorithm variants change the variance of daily sales for small sellers, because exposure volatility can impose operational costs even if average sales increase [30].

An additional macro implication concerns resilience and dependence on a small set of products. Higher concentration can increase platform exposure to supply shocks if dominant products face stockouts or quality issues. The paper therefore computes a concentration-adjusted stockout risk proxy by interacting sales shares with observed stockout frequencies. This is not a structural reliability model, but it quantifies whether algorithm variants allocate demand toward items with more stable availability, which can improve consumer experience even if concentration rises. This illustrates that concentration is not inherently undesirable; its welfare interpretation depends on quality, reliability, and competitive dynamics.

7. Robustness, Limitations, and Policy Discussion

Robustness checks address identification threats that are common in algorithmic settings [31]. For randomized experiments, balance tests confirm that pre-treatment outcomes and key covariates are similar across variants, and clustering accounts for repeated observations per user. Spillovers are examined by testing for changes in aggregate outcomes among users not assigned to treatment during the experiment window, which could occur if sellers adjust prices platform-wide. For staggered rollouts, event-study plots evaluate pre-trends and reveal whether adoption coincides with unusual shocks. Sensitivity analyses exclude major promotional periods and control for marketing campaigns and holidays. For regression discontinuities, density tests on the running variable and continuity checks on covariates support the absence of manipulation near cutoffs, and bandwidth sensitivity assesses stability [32].

Measurement robustness is important because exposure is not equivalent to attention. The analysis uses alternative exposure definitions, including viewport-based impressions, time-in-view thresholds, and scroll-adjusted exposures. Results are similar in sign when using these measures, though magnitudes differ, consistent with the idea that not all impressions are equally salient. Similarly, discovery depends on the lookback window for "previously seen" status, so the paper reports results for multiple windows. The qualitative conclusion that exploration increases discovery is robust, while the estimated conversion tradeoff varies with window choice because longer windows classify more items as previously seen and therefore compress the measured discovery margin.

Interference and displacement pose interpretational challenges [33]. Even with user-level randomization, the treatment changes the composition of what is shown, so product-level effects reflect both own exposure changes and competitive displacement. The analysis addresses this by focusing on group-level exposure shares and within-category concentration outcomes, which are naturally defined under competition for slots. Still, general equilibrium beyond the user can occur through supply-side responses, particularly for advertising and pricing. The paper partially addresses this by measuring changes in sponsored placements and by controlling for contemporaneous advertising intensity, but it does not fully model strategic seller behavior. As a result, medium-run estimates should be interpreted as conditional on the platform's broader environment during the study period [34].

A further limitation is that structural welfare proxies rely on parametric assumptions about choice and consideration. While reduced-form causal effects on observable outcomes are less assumption-dependent, they cannot directly quantify surplus. The structural layer is therefore used primarily for counterfactual accounting and mechanism interpretation rather than for definitive welfare claims. Where possible, the paper triangulates with nonparametric evidence, such as revealed substitution patterns inferred from click sequences and cart edits, to support the direction of mechanism channels.

The policy discussion focuses on levers that platforms can adjust if concentration outcomes are a concern. Diversity constraints in top ranks, calibrated exploration, and constraints on repeated exposure to the same items can broaden exposure, but they may reduce short-run conversion if not carefully tuned [35]. Conversely, objectives that heavily weight conversion and operational metrics can improve reliability and reduce returns but may concentrate demand among operationally advantaged sellers. The analysis suggests that design choices can be framed as tradeoffs among match quality, exploration, and competitive dispersion, and that the relevant objective depends on whether the platform prioritizes short-run efficiency, long-run variety, or seller diversity. The empirical framework supports auditing: by estimating the marginal effects of exposure reallocations on concentration metrics, platforms can quantify how much dispersion is gained per unit of conversion loss, if any, in a particular category.

Finally, the discussion emphasizes that concentration effects are context-specific. In categories with strong quality gradients and high consumer consensus, concentrating exposure on best-performing products may align with consumer welfare proxies, while in categories with heterogeneous tastes and high differentiation, exploration and diversity can improve match quality and reduce reliance on a few incumbents [36]. The paper therefore advocates category-sensitive evaluation rather than a single platform-wide conclusion. It also underscores that transparency about recommendation objectives and constraints can help interpret observed concentration patterns without attributing them solely to algorithmic bias.

8. Conclusion

This paper develops a causal approach to assessing how recommendation algorithms shape product discovery, demand allocation, and concentration in platform-mediated retail. By treating recommendations as an exposure allocation policy under attention constraints, the framework clarifies why naive correlations between being recommended and selling well are uninformative, and it motivates designs that leverage experimental variation, rollouts, and discontinuities to identify causal effects. The empirical analysis decomposes impacts along the funnel from exposure to clicks to purchases, showing how algorithms can increase discovery while producing heterogeneous conversion responses. Demand allocation effects are largely mediated by substitution and displacement within constrained recommendation slots, with early-funnel surfaces exhibiting more scope for expanding engagement than later-funnel surfaces [37].

At the aggregate level, concentration changes are traced back to exposure reallocation and conversion heterogeneity, and counterfactual accounting connects micro-level shifts to seller and product concentration metrics. The results indicate that concentration is sensitive to how top ranks are allocated and to dynamic feedback in learning-to-rank systems, where short-run exposure changes can compound through model retraining. The analysis also highlights that seller-level concentration can move differently from product-level concentration when operational features enter ranking objectives, suggesting that multiple aggregation levels are needed for monitoring.

The paper's contribution is a disciplined empirical and conceptual toolkit rather than a single directional claim about recommendations. Recommendation policies can broaden discovery and diversify demand in some contexts and concentrate demand in others, depending on exploration, personalization, and the alignment of ranking signals with consumer heterogeneity. The framework supports ongoing evaluation and auditing by translating algorithmic changes into interpretable causal quantities and by decomposing concentration effects into actionable channels [38].

References

[1] S. Robbani and I. K. Ningrum, "Non fungible token sebagai aset digital dalam pandangan fiqh muamalah," *At-Tuhfah*, vol. 11, pp. 1–23, 12 2022.

[2] E. Cortelli, "Market monitoring and the detection of market abuses: Fintech frontiers," 10 2018.

[3] J. Morgan and M. R. Baye, *The Economics of E-Commerce*. 10 2016.

[4] N. V. Patel, "Selection bias in two-sided e-commerce marketplaces: A framework for propensity score matching implementation," *Journal of Business Intelligence Systems and Computational Social Science Applications*, vol. 11, no. 5, pp. 1–20, 2021.

[5] E. von Essen and J. Karlsson, "The effect of competition on discrimination in online markets-anonymity and selection.," *PloS one*, vol. 14, pp. 1–18, 8 2019.

[6] J. Janisch and A. Vossen, "How product novelty shapes sales performance of new ventures with deviant business models," *Academy of Management Proceedings*, vol. 2019, pp. 17686–, 8 2019.

[7] R. D. Chikhalkar, "Measurement of reliability as on-line purchase dimension for retail internet sites amongst engineers," *The GSTF Journal on Business Review*, vol. 2, 10 2012.

[8] E.-A. Jung and H. Sung, "The influence of the middle east respiratory syndrome outbreak on online and offline markets for retail sales," *Sustainability*, vol. 9, pp. 411–, 3 2017.

[9] J. Debora, "Impact of e-tail brand experience to brand trust, brand loyalty, and gender as the moderating variable in zalora," 1 2019.

[10] R. Mayya, S. Ye, S. Viswanathan, and R. Agarwal, "Who forgoes screening in online markets and why? evidence from airbnb," *Management Information Systems Quarterly*, pp. 1745–1776, 6 2021.

[11] S. Rahimpour and M. Khabbazian, "Hashcashed reputation with application in designing watchtowers," 1 2020.

[12] W. Li, D. Wu, and H. Xu, "Reputation in china's online auction market: Evidence from taobao.com," *Frontiers of Business Research in China*, vol. 2, pp. 323–338, 7 2008.

[13] *AMCIS - The Effect of Marketer-Generated Content and User-Generated Content on Perceived Product Quality*, 5 2017.

[14] M.-C. Ding, K.-W. Wu, and S.-W. Liu, “Determinants of b2c ec success on market performance of different sizes of firms in taiwan’s e-brokerage sector,” *Journal of Small Business Strategy*, vol. 19, pp. 17–36, 5 2008.

[15] A. Sasso, M. Hernández-Alava, J. Holmes, M. Field, C. Angus, and P. Meier, “Strategies to cut down drinking, alcohol consumption, and usual drinking frequency: Evidence from a british online market research survey.,” *Social science & medicine (1982)*, vol. 310, pp. 115280–115280, 8 2022.

[16] H.-K. Ji, “Characteristics of impulse buying according to price attitude towards internet apparel purchases -focusing on the differences by gender and age-,” *Journal of the Korean Society of Clothing and Textiles*, vol. 37, pp. 737–749, 8 2013.

[17] Y. Gao, C. Kroer, and A. Peysakhovich, “Online market equilibrium with application to fair division,” 6 2021.

[18] G. E. Bolton, C. Loebbecke, and A. Ockenfels, “How social reputation networks interact with competition in anonymous online trading: An experimental study,” 8 2007.

[19] N. V. Patel, “Applying synthetic control methods to address causal identification challenges in the ride-hailing industry,” *Journal of Data Science, Predictive Analytics, and Big Data Applications*, vol. 8, no. 7, pp. 27–49, 2023.

[20] J. Wang, “New normal of online learning: Applying customer satisfaction of learning experiences to obtain competitive advantages on the online market,” 9 2020.

[21] H. D. Wijnholds and M. W. Little, *Regulatory and Marketing Challenges Between and U.S. and EU for Online Markets*. IGI Global, 1 2011.

[22] A. Klobuchar, *Profit Over People*, pp. 301–314. Oxford University PressNew York, 8 2022.

[23] D. P. Mehta, T. M. Kouri, and I. Polycarpou, “Iticse - forming project groups while learning about matching and network flows in algorithms,” in *Proceedings of the 17th ACM annual conference on Innovation and technology in computer science education*, pp. 40–45, ACM, 7 2012.

[24] D. Laffey, K. Plangger, and D. Nel, *Marketing Golden Bytes: A revised online value creation model*, pp. 569–569. Springer International Publishing, 10 2014.

[25] E. S. Soegoto and R. Akbar, “Effect of the internet in improving business transactions with online market methods,” *IOP Conference Series: Materials Science and Engineering*, vol. 407, pp. 012051–, 9 2018.

[26] J. Bozic and F. Wotawa, *ICTSS - Security Testing for Chatbots*, pp. 33–38. Germany: Springer International Publishing, 9 2018.

[27] W. Zou, J. Wang, and J. Yan, “Online markets and trust,” 11 2020.

[28] C. Zhang, F. Wu, and X. Bei, *PRICAI (1) - An Efficient Auction with Variable Reserve Prices for Ridesourcing*, pp. 361–374. Germany: Springer International Publishing, 7 2018.

[29] M. Laouenan, morgane laouenan, P. Deschamps, G. Chapelle, and X. Lambin, “Discrimination on online markets: Evidence from a field experiment,” 3 2021.

[30] H. First and S. W. Waller, “Internet markets and algorithmic competition: The rest of the story,” 8 2017.

[31] Y. A. B. El-Ebiary, A. Hatamleh, S. Saat, K. T. Amayreh, R. Karim, S. Bamansoor, and M. H. Yusoff, “Online market between problems and challenges,” 4 2021.

[32] L. Janezic, “Personalization friend or foe,” 7 2014.

[33] N. Duch-Browne, L. Grzybowski, A. Romahn, and F. Verboven, “Are online markets more integrated than traditional markets? evidence from consumer electronics,” 8 2020.

[34] F. van Beetz, “An exploratory research on online market learning methods,” 8 2016.

[35] N. V. Patel, “Adaptive experimentation in marketplaces: Balancing exploration and revenue optimization,” *Frontiers in Health Informatics*, vol. 11, pp. 760–786, 2022.

[36] A. Agrawal, C. Catalini, and A. Goldfarb, “Crowdfunding: Geography, social networks, and the timing of investment decisions,” *Journal of Economics & Management Strategy*, vol. 24, pp. 253–274, 5 2015.

[37] A. Nair and R. Mathews, *Challenges and Solutions in Recommender Systems*, pp. 188–194. Springer International Publishing, 3 2020.

[38] A. R. Reuber and E. Fischer, “Signalling reputation in international online markets,” *Strategic Entrepreneurship Journal*, vol. 3, pp. 369–386, 12 2009.