Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Variety Seeking in High-Frequency Consumption: Implications for Targeted marketing

Fan Zhang

fan.zhang@novasbe.pt

Universidade Nova de Lisboa

Carol Hengheng Lu

Amazon (United States)

Zhijie Lin

Tsinghua University

Research Article

Keywords: Variety Seeking, Targeted Marketing, Online Platform, State Dependence

Posted Date: May 20th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9407887/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Variety Seeking in High-Frequency Consumption: Implications for Targeted Marketing

April 14, 2026

Abstract

This paper studies consumers' variety-seeking preferences and their implications for targeted marketing on a food delivery platform. Our analysis reveals that a substantial fraction of consumers exhibit variety-seeking preferences, with an average consumer willing to pay 23.1% less to reorder from the same seller they purchased from in the previous period. Through counterfactual analysis, we find that optimizing rankings to account for variety-seeking preferences can enhance revenue, consumer welfare, and purchase probability. Furthermore, we find that the optimal targeted pricing strategy involves a general price increase with a discount for recent consumers, and consumers' variety-seeking preferences soften price competition.

Keywords: Variety Seeking, Targeted Marketing, Online Platform, State Dependence

1 Introduction

Targeted marketing has been a widespread business practice because of its unprecedented ability to engage customers on a personal level. At its core, the effectiveness of the targeted marketing strategy hinges on how well it can recognize and respond to the unique needs and interests of each consumer. Rossi et al. (1996) find that even a short purchase history of consumers can improve profitability more than demographic information, underscoring the importance of leveraging consumer behavioral data to identify the unobserved persistent consumer heterogeneity.[1] Besides the static persistent preference, individual purchase history is also informative about

[1]Since then, using behavioral data to identify the unobserved persistent preference heterogeneity that is not captured by demographic information has been the focus of follow-up research on personalized marketing. For example, Besanko et al. (2003); Pancras and Sudhir (2007); Howell et al. (2016); Morozov et al. (2021); Donnelly et al. (2021); Smith et al. (2023) study this on the topic of personalized pricing.

temporary dynamics in demand. When the persistent preference and temporary dynamics point in opposite directions, failing to account for the dynamics will decrease consumer welfare through distraction and satiation, preventing the optimal use of the marketing resources. Accounting for the temporary negative demand has important implications for targeted marketing practices, including pricing, ranking, recommendation systems, and retargeting advertising.[2] Compared to the optimal targeted marketing practices, consumers are exposed to and nudged toward their favorite products too soon, too much, and too frequently.

In this paper, we empirically estimate the relative magnitude of the persistent preference and temporary dynamics in demand. We highlight consumers' variety-seeking preferences as an important source of demand dynamics that can be estimated from purchase history data and utilized for targeted marketing. We document a substantial fraction of consumers with variety-seeking preferences and study the implications of consumers' variety-seeking preferences for targeted ranking and competitive targeted pricing.

We use proprietary data from a peer-to-peer (P2P) homemade food delivery platform, which provides several advantages for examining variety-seeking preferences and their implications for targeted marketing. First, food is consumed frequently with clearly defined daily demand, which eliminates the typical issues of long purchase cycles, durability, and stockpiling with consumer packaged goods. Second, since the platform offers uniquely available products, the outside options exclude consumption of these specific products, unlike traditional packaged goods markets where unobserved channels may exist. Moreover, most of the orders are workday lunch orders, reducing potential confounding from intrahousehold heterogeneity (Che et al., 2003). Finally, we observe marketing variables, including pricing and ranking, which enable us to estimate the effect of these marketing instruments and conduct counterfactual analysis on these instruments while considering consumers' variety-seeking preferences.

We first report descriptive evidence on consumers' variety-seeking preferences. In our data, the probability that a consumer switches to a different seller from their previous seller choice is 81%, which stands in stark contrast to the switching probability found in the existing literature on state dependence. To better account for alternative driving forces of switching probability besides variety-seeking preferences, such as choice set size and heterogeneous preferences, we

[2]See Appendix E.2 for examples of personalized recommendations that push recommendations on favorite sellers ordered very recently. Such practices emphasize persistent heterogeneous preferences, but have the hazard of potentially neglecting the dynamics in demand driven by state-dependent preferences. Another example of inefficient targeted advertising due to neglecting negative state dependence is retargeting: For instance, `https://www.digitaltrends.com/features/wh-do-you-see-ads-stuff-already-bought/` illustrates a case where a recent consumption is followed by a decrease in demand, and ignoring this in retargeting advertising reduces both profits and consumer welfare.

use a permutation-based test and show that consumers in our data switch across different kitchens more frequently than in a random reshuffle benchmark, where the choice sequences are randomly reshuffled with market shares fixed. This provides strong suggestive evidence of consumers' variety-seeking preferences.

We then model consumers' discrete demand for kitchens, allowing for negative state dependence. We estimate a random coefficient multinomial logit model to measure variety-seeking preferences, controlling for heterogeneous preferences, product attributes, and marketing variables. We find that, on average, consumers exhibit variety-seeking preferences and are willing to pay 23.1% less to reorder from the same seller, despite persistent preference. We find rich heterogeneity in consumers' state-dependent preferences: 70.5% of consumers have variety-seeking preferences, while 29.5% of consumers exhibit inertia. We rule out potential confound of spurious state dependence (Heckman et al., 1981), and alternative sources of structural state dependence that result from information gained from consumption experience instead of changes in preference, including learning and new product discovery.

We use our demand estimates to study the implications of variety-seeking preferences for targeted marketing. For targeted ranking, we analyze profits, welfare, and purchase probability under various counterfactual ranking schemes. We find that optimizing ranking algorithms with variety-seeking preferences comprises 14.6% of the revenue improvement, 17.2% of the consumer welfare improvement, and 18.3% of the purchase probability improvement out of the total ranking effect. Most of the welfare gain comes from consumers purchasing a more desirable seller who is ranked lower under the suboptimal ranking, such that the consumers choose the outside option or a worse match ranked at a higher position.

We next use our demand estimates to study how variety-seeking preferences affect competitive targeted pricing. Existing theoretical studies (Seetharaman and Che, 2009; Sajeesh and Raju, 2010; Zeithammer and Thomadsen, 2013) generate mixed predictions on the effect of variety-seeking preference on price competition and suggest that it depends on the magnitude of differentiation in persistent preference and variety-seeking preference, which makes this an empirical question.

We start by studying how optimal targeted pricing strategies affect equilibrium prices, profits, and consumer welfare at the estimated variety-seeking level. We find that optimal targeted pricing results in lower prices to own consumers and higher prices to rival kitchens' consumers, which implies a price increase with quantity discount, compared with the prices ignoring the variety-seeking effect. We further investigate how variety-seeking preferences would affect the price competition level. We demonstrate the directions of pressures on prices resulting from

the sellers' current period and future profit incentives when facing variety-seeking consumers. We find that variety-seeking preferences soften price competition and lead to a higher average equilibrium price compared to the no state dependence case. Moreover, the prices to a kitchen's own consumers decrease with the variety-seeking level while the prices to rival kitchens' consumers increase with the variety-seeking level. These findings have important implications for competitive targeted pricing and suggest that sellers could increase profits by offering their own consumers coupons for quantity discounts. However, under current common industry practice, popular food delivery platforms do not support this type of state-based targeted pricing.[3]

This paper contributes to several strands of literature. Existing research on consumer state-dependent preferences has focused on data from consumer packaged goods(CPG), where a vast majority of existing evidence has been on inertia (Dubé et al., 2010; Osborne, 2011; MacKay and Remer, 2019; Kong et al., 2022; Levine and Seiler, 2023). Studies on variety seeking, pioneered by McAlister (1982) [4], mostly use stochastic choice models (Givon, 1984; Kahn and Raju, 1991; Trivedi et al., 1994; Chintagunta, 1998), or a hybrid version of stochastic choice models and utility-based models (Bawa, 1990; Seetharaman and Chintagunta, 1998; Chintagunta, 1999; Park and Gupta, 2011; Smith et al., 2020). We contribute to this stream of literature by documenting variety-seeking preferences in the novel context of online platforms. Compared with stochastic choice models, our utility-based choice model better accommodates the variation in the choice environment and enables us to conduct counterfactual analysis of alternative marketing practices in personalized ranking and pricing.

Our work also relates to the literature on consumer preference for variety. Existing studies have different definitions of variety, such as static variety measured by the size and diversity of provided assortments and purchased product sets (Van Herpen and Pieters, 2002; Brynjolfsson et al., 2003; Broda and Weinstein, 2006; Quan and Williams, 2018; Datta et al., 2018; Ershov, 2020; Holtz et al., 2020; Natan, 2025; Trindade and Caldieraro, 2025). For the static variety, the order of the choice sequence does not matter, and consumers' choices are aggregated along the time dimension. The definition of variety seeking in our project is negative state dependence: a disutility from consuming the same product consecutively. We study consumers' preferences for dynamic variety over time where the order of the choice sequences matters.

With the development of digital platforms and increasing availability of individual-level

[3]Section E.1 provides more details about the current targeted pricing industry practice.

[4]McAlister and Pessemier (1982) provides a review of the different psychological factors behind the variety-seeking behavior, such as satiation, boredom, experience seeking, or thrill and adventure seeking. In this project, we do not try to differentiate these psychological mechanisms; instead, we summarize them by negative state dependence preference.

ranking data, studies on improving personalized ranking algorithms have attracted the attention of marketing researchers (Jeziorski and Segal, 2015; Yoganarasimhan, 2020; Derakhshan et al., 2022; Compiani et al., 2024; Abaluck et al., 2025). We contribute to this stream of literature by examining the role of state-dependent demand dynamics in the optimization of ranking algorithms. We highlight consumers' variety-seeking preferences as an important factor in personalized ranking algorithms to improve consumer welfare, purchase probability, and platform revenue.

Finally, we contribute to a broader literature that studies state dependence and pricing. Che et al. (2003) studies the forward-looking pricing behavior in markets with state dependence within a two-period model. Dubé et al. (2008) and Dubé et al. (2009) study the pricing implications of switching costs with infinite horizon, under monopoly and oligopoly market structures, respectively.[5] Zhang and Krishnamurthi (2004) and Zhang and Wedel (2009) study the implications of consumers' state-dependent preferences for targeted pricing in a three-period off-equilibrium analysis, where the target brand can conduct targeted pricing while keeping other brands' prices fixed. We contribute to this literature by studying price competition equilibrium with variety-seeking consumers, allowing the sellers to compete with state-based targeted pricing. Our study sheds light on the optimal strategy and provides empirical evidence on the theoretically ambiguous questions of how negative state dependence affects price competition.

The rest of the paper proceeds as follows: Section 2 describes the background and data, Section 3 provides descriptive evidence of variety-seeking, Section 4 presents and estimates the demand model. Section 5 reports results from counterfactual analysis of targeted ranking and targeted pricing, and Section 6 concludes.

2 Background and Data

We study data from a P2P online food ordering platform. The platform, the first P2P homemade food delivery platform in China, was founded in Beijing in 2014, connecting households offering home-cooked food with nearby consumers. Until 2017 there were more than a million buyers who ordered food from sellers via the mobile app, and twelve thousand sellers who were mostly homemakers and the retired.[6]

The data include detailed records of all orders placed by buyers on the platform in Beijing

[5] Farrell and Klemperer (2007) and Villas-Boas (2015) provide theoretical analysis of how switching costs affect price competition.

[6] Appendix A provides more background information about the market structure of the food delivery industry, consumers' daily and weekly order trends, app display, and kitchen cuisine distribution.

from November 2016 to January 2017. For each order, we observe the identity and location of the buyer and seller, time of order, order fee, dishes ordered, delivery fee, coupon offered, and rating and review for the order. We also observe the ranking of kitchens shown on consumers' homepages and the final purchase outcome. We focus on the frequent users who ordered more than 10 times from the platform during the three-month window.[7] Because food consumption happens frequently on a daily basis, a natural definition of one time period is one calendar day. We define outside options as the days when consumers opened the platform and didn't eventually purchase; in those cases, we observe the choice sets in their rank lists. These periods when consumers used the app and checked out the kitchen list are important even though consumers did not order from the platform, since the pricing and ranking strategies in these periods may have a substantive effect on the extensive margin. [8]

The final sample contains 6,629 consumers, 6,664 kitchens, and 125,646 orders. One consumer on average placed 18.95 orders during the three-month period, and ordered 2.91 times a week, conditional on ordering that week, mostly on weekdays and for lunch. There are on average 14 kitchens available nearby for each consumer per day, and an average consumer ordered from 9.28 distinct sellers throughout the three-month period. Detailed summary statistics are shown in Table 1.

[7] These frequent users take 10% of the consumers and 50% of the orders and revenues.

[8] We dropped the periods when consumers didn't use the platform. This definition of outside option implies that our estimate of variety seeking and the corresponding counterfactual analysis of extensive margin are conservative. For example, suppose $\{A, B\}$ represent different kitchens on the platform, O represents the outside option, then the following order sequence presents consecutive orders of the same kitchen A, whereas the orders were not immediately after each other.

$$ABA \quad \underbrace{OOOOO}_{\text{no platform usage}} \quad A \longrightarrow ABAA$$

.

Table 1: Descriptive Statistics

	Mean	Std. Dev.	10 Pctile	Median	90 Pctile
Consumer-Level Data					
No. Platform Uses	40.03	15.22	22	38	62
No. Orders	18.95	8.20	13	16	30
No. Orders per Week	1.91	0.80	1.09	1.70	3.01
No. Orders per Week Conditional on Ordering	2.91	0.94	2	3	4
Prob. of Week with Ordering	0.83	0.13	0.67	0.83	1
No. Kitchens (Daily Average)	14.04	5.70	6	15	31
Distinct Kitchens Ordered	9.28	5.21	3	7	15
Distinct Cuisines Ordered	3.65	1.91	1	3	6
Distinct Dishes Ordered (Per Kitchen Ordered)					
Kitchen-Level Data					
Rating (1-5)	4.89	0.15	4.49	4.87	5
Menu Size	25.32	9.81	3	23	49
Average Dish Price (CNY)	25.21	9.45	14.29	23.32	29.15
Average Price of a Basket of Food (CNY)	50.23	17.32	26.19	50.86	58.49
No. Orders per Week	18.76	7.64	1.72	10.40	48.25
Delivery Radius (km)	2.27	1.31	2	3	9
No. Days Available per Week	3.96	2.06	1	4.23	6.46
No. Appearance in Choice Sets Daily	17.61	15.96	3.67	13.07	37.01
Order-Level Data					
Order Payment (CNY)	45.46	22.37	22	48	68
Delivery Fee (CNY)	2.89	0.77	0	3	3
Distance (km)	1.68	0.99	1.54	2.43	2.76
No. Dishes per Order	2.38	1.01	2	3	5
No. Consumers	6,629				
No. Sellers	6,664				
No. Orders	125,646				
No. Platform Uses	265,375				

Note: The table presents summary statistics of our sample. We focus on frequent users who ordered more than 10 times during the sample period of November 2016 to January 2017.

3 Descriptive Evidence

In this section, we provide motivating evidence of consumers' variety-seeking preferences. As variety seeking is about the demand dynamics across different time periods, we study the switching probability as a measure of consumer choice variety. We find that in our data consumers stay with the same kitchen they chose in the last period only 19% of the time, for 36% of the time they switched to a different kitchen on the platform, and for 45% of the time they switched off the platform.[9]

The switching probability of 81% in our data stands in stark contrast to that found in the existing literature, which normally finds a repurchase probability of 80-90% (Dubé et al., 2010). However, the switching probability is also affected by the size of choice sets and the relative magnitude of consumers' persistence preferences. A switching probability of 81% has different implications when consumers face only two options each period, and the case when consumers face hundreds of options. Switching away from a favorite kitchen is also more informative than switching away from a less preferred kitchen. Thus, in order to have a clearer evaluation of the switching frequency, we construct a permutation-based benchmark to compare the switching frequency we observed in the data with that when the order of choices does not matter, given the observed market shares.[10]

To account for consumers' persistent kitchen preference, we construct a permutation-based switching probability benchmark by randomly shuffling the order sequence of each consumer 100 times while holding fixed the market shares for each kitchen ordered. Table 2 presents an example of how the permutations are implemented. We compare the switching probability observed in the data and the average switching frequency in the randomly shuffled benchmark. The distribution is shown in Figure 1. Panel (A) presents the scatter plot with each dot representing one consumer. For consumers with actual switching probability greater (less) than the randomly reshuffled benchmark level, they switch more (less) frequently in the data than the case when the order of choices does not matter, and exhibit variety-seeking preferences (inertia).

[9] As the products on our platform are homemade food only available through the platform channel, instead of standardized branded food, it is reasonable to assume that people consumed different products from the ones on the platform when they chose the outside option.

[10] Appendix B presents details of the statistical test. A similar permutation-based test is used by Miller and Sanjurjo (2018). By permutations, we preserve the randomness in choices and the number of orders for each chosen kitchen in the choice sequences while removing the state dependence. If the observed switching probability is not significantly higher than the random permutation switching probability, then we can't reject the null hypothesis that consumers do not have state-dependent preferences. Such tests of checking for an association that should be present if the design is flawed but not otherwise are called placebo tests by some existing literature (Larsen et al., 2021). See Eggers et al. (2021) for a review of placebo tests.

Table 2: Example of Permutation Test

Observed Choice Sequence	Switching Probability	Permutations	Switching Probability
		CBAAA	0.5
		ABAAC	0.75
ABACA	1	BCAAA	0.5
		ACABA	1
		CABAA	0.75
		...	...

Note: The table illustrates how the permutation test is implemented with an example. Given an observed choice sequence in the data, we hold fixed the market shares and randomly reshuffle the order of choices. The distribution of switching probability from the permutations recovers the distribution of switching probability under the null hypothesis that there is no state-dependent utility and the order of choices does not matter.

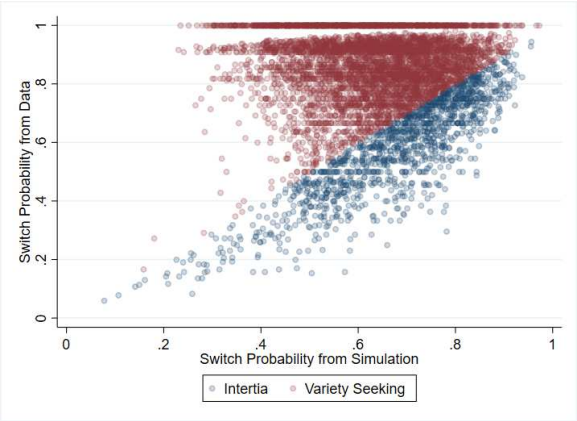
A majority of consumers exhibit variety-seeking preferences (79%), and a minority of consumers have inertia (21%). The figure also suggests that there is heterogeneity in consumers' state dependence preferences, and we formally model the heterogeneity through a random coefficient model in the structural estimation. Panel (B) presents the cumulative distribution function of the switching probability in data, in the random reshuffled benchmark, and the 95% confidence bands of the reshuffled sequences. For most parts of the distribution, the observed switch probability lies well outside the 95% confidence bands of the random reshuffled distribution, suggesting that the switch probability observed in the data is not simply from randomness in choices. The mean switch probability in the random reshuffled benchmark is 64%, with a 95% confidence interval of [47%, 79%], which suggests that the consumers switch 17% more than a simple random choice. The average switching probability across consumers indicates variety seeking with a significance level $p < 0.05$.[11]

The analysis above is based on stochastic choice models and does not account for price, availability, changing of the choice set, and other product attributes such as cuisine type and menu size, or time-related effects such as day of the week, week of the month, and holiday effects. To formally account for these confounding factors, to measure consumers' state-dependent preferences, and to further conduct counterfactual predictions, we need structural modeling.

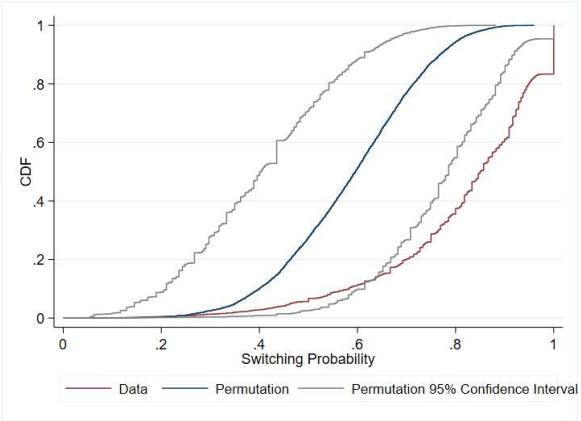
[11]We conduct this pooled test with the average of standardized switching probability $\bar{p} = \frac{1}{N}\sum_i \tilde{p}^i$, where, for each consumer, the switching probability $\tilde{p}^i$ is normalized by shifting its mean and scaling its standard deviation under H_0. In this case: $H_0 : \mathbb{P}(y_t \neq s_t | s_t \neq 0) = p^i$ for all t, i, where y_t denotes current period choice and s_t denotes previous period choice. We stratify the permutations by consumer to recover the sampling distribution of average switching probability under null hypothesis. Moreover, a Kolmogorov-Smirnov test comparing the actual CDF of switching probability to the CDF of the mean switching probability from permutations yields a pj0.05.

Figure 1: Descriptive Evidence

(A) Switching Probability Distribution



(B) Switching Probability CDF



Note: Panel (A) shows the switching frequency observed in the data relative to that in the randomly reshuffled benchmark, where the order of choices does not matter given the observed market share. One dot represents one consumer. The 45-degree line represents the cases when consumers switch as frequently as the randomly reshuffled benchmark. The consumers above the 45-degree line switch more frequently in the data than the random reshuffle benchmark and exhibit variety-seeking preferences, whereas the consumers below the 45-degree line switch less frequently than the randomly shuffled benchmark and exhibit inertia. Panel (B) presents the cumulative distribution function of the switching probability in data, in the random reshuffled benchmark, and the 95% confidence bands of the reshuffled sequences.

4 Demand Estimation

This section provides our demand model, parametric assumptions, estimation strategies, and identification inference.

4.1 Model

We introduce the structural model in this section. We assume that consumers maximize their utility given the available choice set, which we observe from the search result data.[12] The indirect utility that consumer i derives from kitchen j on platform in period t is

$$u_{ijt} = \beta_{ij} - \alpha_i p_{jt} + \eta r_{ijt} + \gamma_i \mathbb{1}(s_{it} = j) + \psi \cdot X_{ijt} + \epsilon_{ijt}. \quad (1)$$

The utility from choosing the outside option is $u_{i0t} = \epsilon_{i0t}$.

We use $\beta_{ij} \sim N(\mu_j, \sigma_j^2)$ to capture the match value of consumer i and kitchen j. The mean μ_j

[12] We keep the periods when consumers ordered from kitchens on the first page of search results, which account for 92% of all orders. We drop the purchases on pages beyond the first page because of the potential selection problem: consumers who search beyond the first page are more likely to be searching for a specific niche kitchen, and the selection effect will bias the estimates of position cost.

captures the unobserved quality of the kitchen, and the variance allows for heterogeneity across different consumers' tastes, which is centered around the mean quality. Different kitchens are allowed to have different mean quality and dispersion of tastes. Price sensitivity is captured by $\alpha_i \sim N(\mu_\alpha, \sigma_\alpha^2)$, we use random coefficients to allow for heterogeneity across consumers. Since our main focus is on the state dependence preference, we use a position cost to capture the ranking effect instead of imposing a full structural search model: r_{ijt} is the rank of kitchen j in consumer i's rank list in period t, and η represents the position cost from choosing a lower-ranked kitchen.[13]

In each period t, consumer i has chosen one option in the last period, captured by state variable $s_{it} \in \mathcal{J}_{it}$. The choice set $\mathcal{J}_{it}$ includes the outside option (represented by option index zero) and the observed kitchen list for consumer i in period t. If the consumer chose option j in the last period, $s_{it} = j$, and subsequently purchases product, $k \in \mathcal{J}_{it}$, then his or her state becomes, $s_{i(t+1)} = k$.[14]

The coefficient $\gamma_i \sim N(\mu_\gamma, \sigma_\gamma^2)$ captures the state dependence when consumer i chooses the same kitchen on the platform. A positive state dependence coefficient suggests inertia, and a negative state dependence coefficient corresponds to variety seeking. The covariate vector X_{ijt} includes the menu size, cuisine, day of the week, week of the month, and holiday fixed effects. The idiosyncratic taste shock ϵ_{ijt} is drawn from Type I Extreme Value distribution.

4.2 Estimation

We use the Simulated Maximum Likelihood method to estimate the model. We assume that ϵ_{ijt} is distributed Type I Extreme Value and is independently and identically distributed (i.i.d) across consumers, time, and kitchens. U_{ijt} represents the deterministic part of the utility, i.e., $u_{ijt} = U_{ijt} + \epsilon_{ijt}$. Then the probability of consumer i choosing option j is

[13] See Ansari and Mela (2003),Ghose et al. (2014), and De los Santos and Koulayev (2017) for examples of a similar method to capture the ranking position effect. Partial information models in which previous period choice has a higher probability of entering the consideration set, or a lower search cost, will imply that our estimate of variety seeking is conservative, since the information frictions work in favor of the previous period choice because memory and salience of the recent choice are strongest among all options. Appendix C.5 reports results on the out-of-sample fit of this model specification.

[14] As we study high-frequency food consumption, our definition of the state variable when consumers choose the outside option is different from common practice in the literature, such as Dubé et al. (2010), which often assumes that the state variable does not change when the consumer chooses the outside option. In the context of our data, it is more reasonable to assume that consumers have lunch every day and have demand for the platform on the days when they open the platform, and that the outside option is different from the kitchens on the platform. As our platform was the only P2P platform at the time, the products were exclusively available through the platform. This is an advantage compared with platforms where the sellers are restaurants that can be multi-homing, such as Uber Eats and DoorDash.

$$\mathrm{P}_{ij} = \frac{\exp\left(U_{ijt}\right)}{\sum_{k \in \mathcal{J}_{it}} \exp\left(U_{ikt}\right)}. \tag{2}$$

The likelihood function for a sequence of consumer i' purchase decisions $\mathbf{y}_\mathbf{i}^*$ is

$$\mathcal{L}_i\left(\mathbf{y}_\mathbf{i}^* \mid \theta\right) = \frac{1}{R} \sum_{r=1}^{R} \left(\Pi_{t \in T} \Pi_{j \in \mathcal{J}_{it}} \left(\frac{\exp\left(U_{ijt}^r\right)}{\sum_{k \in \mathcal{J}_{it}} \exp\left(U_{ikt}^r\right)} \right)^{y_{ijt}^*} \right), \tag{3}$$

where r is the index for draws of random coefficients. We use 100 draws in the simulations.

4.3 Identification

The main challenge of evaluating past choices' effect on current choices is to disentangle consumers' heterogeneous preferences and structural state dependence. To solve this problem, we first control for the unobserved heterogeneity by adding the random coefficients β_{ij}, which captures both vertical heterogeneity across kitchens by the mean quality and also horizontal taste heterogeneity by the dispersion level of taste distribution.[15] We assume that the unobserved heterogeneity is stable over time, which is a reasonable assumption since we focus on a relatively short time period. Thus the persistent unobserved heterogeneity is well captured by the random coefficient β_{ij}. The identification of the state dependence parameters separately from heterogeneous preferences then comes from the variation of purchase probability right after a previous purchase induced by a rich set of frequent exogenous shocks to availability, prices, and rankings over time in our data.

Availability The special composition of the supply side of the P2P platform provides exogenous shocks to kitchens' availability. Since the kitchen owners on the platform are mostly homemakers and retired workers, they make flexible decisions of availability based on personal life arrangements, such as childcare and travel. The median number of orders per kitchen per week is 10; it is reasonable to assume that the kitchens are not making sophisticated strategic decisions on availability, given the sales volume, potential profit margin, and the lifestyle of kitchen owners.

Price The variation in prices in our data comes from differences across kitchens, different

[15] As there are 6,664 kitchens in our data, it is intractable to allow for a different distribution for each kitchen. To deal with this dimensionality problem, we follow Bonhomme et al. (2019), Bonhomme et al. (2022), and use the k-means method to group all kitchens into 16 clusters based on observed average purchase probability, prices, and kitchen characteristics. Kitchens in each cluster share the same distribution of β_{ij}. The grouping brings approximation error but reduces incidental parameter bias. We restrict the approximation error by following the data-driven rule provided in Bonhomme et al. (2022). The reduction in incidental parameter bias comes from two sources: utilizing more informative moments about the latent types and regularization.

choices of menu items, menu changes, delivery fees, and promotional discounts. The potential endogeneity of time-persistent price variation across different kitchens is captured by β_{ij}. We discuss the price variation over time caused by menu changes, delivery fees, and promotional discounts below. We do not model consumers' dish-level choices. Instead, we observe the kitchen's menus over time and construct the price variable based on the average price of a basket of food at the restaurant. Specifically, we calculate a weighted average of menu prices, with each item price weighted by the aggregate popularity. Sometimes kitchens change their menu by adding or deleting dishes and adjusting prices. When kitchens need to change menus permanently, they need to schedule an in-person check with the platform before the change is implemented. As other active kitchens also need to be checked in person at least every three months, kitchens can not control or predict the timing of menu changes, we assume that the menu changes are exogenous. We control for the menu size to capture the potential diversity variation from menu changes. Beyond the menu prices (tax included), we added the delivery fee. Kitchens set a delivery radius within which they deliver food to consumers and set a delivery fee.[16] Thus, variation in delivery fee over time is generated by variation in consumer's geographic location and distance to the kitchens, and is assumed to be orthogonal to food consumption utility. The platform has platform-level promotions such as Tuesday half-price, free delivery on Friday, and coupons on every 17th day of the month.[17] These promotions are provided by the platform to all consumers and could be potentially correlated with demand. We controlled for day of week, week of month, and holiday fixed effects to deal with the potential endogeneity problem. The frequent temporary platform-level sales, delivery fee variations, and kitchen-level menu adjustments provide exogenous variation in prices after controlling for day of the week, week of the month, and holiday fixed effects, which identify consumers' price sensitivity.

Rank The ranking algorithm in our data is not personalized based on individual consumption history. Instead, the default ranking is based on distance, price, monthly sales, and ratings. The time-persistent unobserved heterogeneity is captured by the random coefficient β_{ij}. The remaining factors that vary over time and could affect rank are availability, new entries, distance, ratings, price, and monthly sales. Since our platform is a P2P platform, some kitchens may be unavailable on certain days, and change the relative ranking positions of the choice set. There

[16]The sellers can choose to deliver by themselves, use the platform's delivery service, or use their own third-party delivery service. In all cases, the consumer needs to pay the delivery fee. If sellers choose to deliver themselves, they do not need to pay extra delivery costs to the platform or a third-party delivery service. We do not observe great variation in delivery fee across orders, and we assume the heterogeneity in delivery fee is driven by supply-side cost heterogeneity, such as the transportation cost of the sellers.

[17]The pronunciation of the number '17' in Chinese sounds similar to 'want to eat'. This is an example of a random price discount that alters prices but not consumers' preferences for food.

can also be some new kitchens entering the market, which could change the ranking positions of existing kitchens. We assume these changes are exogenous due to the uncontrollable scheduling of in-person platform check, and the P2P nature of the platform. Geographic distance between consumers and kitchens varies as consumers change location during the period. We assume that within the delivery radius limit, geographic distance does not enter utility and provides an exogenous shock to rankings over time that identifies the position effect. Monthly sales volume is an aggregate variable that captures the population choice in past periods. The variation in monthly sales over time comes mostly from kitchen availability decisions. We don't include it in the utility function and assume that it is exogenous after controlling for the kitchen fixed effects. The rating variable has very limited variation across kitchens and across time periods. The median rating is 4.87. The median of the ratio of kitchens' overtime rating standard deviation to mean rating is 0.07%. We don't include rating in the utility function as the limited variation makes it less informative given kitchen fixed effects. The frequent temporary kitchen availability changes, new entries, and geographic distance changes provide exogenous variation in ranks after controlling for day of week, week of month, and holiday fixed effects, which identify consumers' position cost. We perform a robustness check on rank endogeneity by adding distance and ratings in the utility function and discuss more details on potential rank endogeneity in Section 4.5.3.

4.4 Results

The estimation results are in Column (1) of Table 3. The price coefficient is negative and significant.[18] The implied position cost is 4.1% of the average payment per order. We find γ_i has a negative mean -0.231 and a standard deviation of 0.430. That suggests that 70.5% of the consumers are variety seeking, with $\gamma_i < 0$; 29.5% of the consumers have inertia, with $\gamma_i > 0$. The heterogeneity level estimated from structural estimation is similar to what we see in the reduced form evidence in Figure 1 panel (B). The ratio of the mean of γ_i to α_i represents consumers' willingness to pay for variety seeking, which is 23.1% of the average payment per order.

[18]The average own price elasticity is 1.3, which is inelastic compared with estimates using data from related industries such as Natan (2020). The difference is likely due to the monopoly nature of the platform. In Natan (2020), there is competition on both the platform and seller level, and buyers and sellers can both be multihoming. Whereas in our setting, the platform is a monopoly in the homemade food subcategory, and the food is exclusively available from this channel, which enables the platform to price at the more inelastic region of the demand curve.

Table 3: Estimation Results.

		(1) Baseline	(2) Learning	(3) New Product	(4) Rank Robustness
State Dependence	μ_γ	-0.231***	-0.216***	-0.247***	-0.293***
–Kitchen level (γ_i)		(0.040)	(0.047)	(0.050)	(0.091)
	σ_γ	0.430***	0.728***	0.697***	0.577***
		(0.054)	(0.060)	(0.055)	(0.075)
Price (α_i)	μ_α	-0.022***	-0.019***	-0.023***	-0.021***
		(0.005)	(0.003)	(0.005)	(0.004)
	σ_α	0.009***	0.011***	0.010***	0.007***
		(0.002)	(0.001)	(0.002)	(0.006)
Rank		-0.041***	-0.032***	-0.038***	-0.034***
		(0.004)	(0.004)	(0.003)	(0.004)
Cumulative Experience	γ^E		-0.003		
			(0.002)		
New Kitchens	γ^N			-0.007	
				(0.027)	
Distance					-0.001*
					(0.001)
Rating					0.001
					(0.002)

Note: The table shows estimation results after controlling for cuisine, day of week, week of month, holiday, menu size, distance, rating, monthly sales, and kitchen fixed effects. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Column (1) is the baseline model. Columns (2)-(4) are robustness checks. Column (2) tests for learning with the model specified in section 4.5.1. Column (3) tests for new product discovery with the model specified in section 4.5.2. Column (4) tests for robustness to the potential rank endogeneity problem.

4.5 Robustness Checks

In this section, we investigate several alternative models that could be sources of the state dependence we identified, including learning and new product discovery. We do not impose specific structural models of learning or new product discovery, which would involve some strong structural assumptions on consumer behavior. Rather, we focus on aspects of consumer behavior that differentiate learning or new product discovery explanations from variety seeking and that can be directly observed in our data. We also discuss and test for additional sources of potential

rank endogeneity.[19]

4.5.1 Learning

One alternative explanation of the negative state dependence we identify is consumer learning. Consumers can switch away from the recent choice not because of the disutility of repurchase, but from extra information or newly resolved uncertainty from consumption, especially at the beginning stage of platform experience (Ackerberg, 2003; Narayanan and Manchanda, 2009; Osborne, 2011). To investigate this possibility, we define consumer-kitchen level consumption experience as the cumulative number of a consumer's orders of a kitchen, E_{ijt}. We interact the state dependence variable with the experience variable to compare the learning and variety-seeking models. Specifically, the indirect utility that consumer i derives from kitchen j on platform in period t is

$$u_{ijt} = \beta_{ij} - \alpha_i p_{jt} + \eta r_{ijt} + \gamma_i \mathbb{1}(s_{it} = j) + \gamma^E \mathbb{1}(s_{it} = j) E_{ijt} + \phi^E E_{ijt} + \psi \cdot X_{ijt} + \epsilon_{ijt}. \quad (4)$$

Under learning, the interaction term should reduce state dependence as brand experience accumulates, i.e., $\gamma^E > 0$. This will generate a convergence pattern in consumers' choice processes, which is prevalent in learning models: they will switch away less from more frequently ordered kitchens, since more information is gained and more uncertainty is resolved, and eventually behave in accordance with a standard choice model with no uncertainty. In contrast, in variety-seeking models, accumulative experience does not affect the state dependence. Table 3 Column (2) shows the estimation results. The coefficient of the interaction term is close to zero and insignificant, which suggests that the variety-seeking preference we identify does not decrease with accumulative experience between consumers and kitchens.

4.5.2 New Product Discovery

In Section 4.5.1, we show that negative state dependence does not decrease with the accumulative experience between consumers and kitchens, which suggests non-convergence in choice behavior and serves as evidence against learning models. However, the non-convergence can also come from constant new entries throughout the time period; consumers could be trying out new options as an investment for discovering a potentially good choice for the future (Ershov, 2020). To examine this possibility we estimate the following specification, in which we add

[19] In Appendix C we provide more robustness discussions about initial conditions bias, inter-purchase time span, variety seeking on the dish level, and spurious state dependence from negative autocorrelation in unobserved shocks. We also report results from a test on the model's out-of-sample fit.

a variable indicating the number of new kitchens in consumer i's choice set in period t, N_{it}, where a new kitchen is defined by the first time the kitchen appears in the consumer's choice set.[20] Under new product discovery, the interaction term should increase state dependence as more new products in the choice set will encourage consumers to switch away from their former purchase, i.e., $\gamma^N < 0$. However, in variety-seeking models, the state dependence should not be affected by the number of new products available in the choice set. Specifically, the indirect utility that consumer i derives from kitchen j on platform in period t is

$$u_{ijt} = \beta_{ij} - \alpha_i p_{jt} + \eta r_{ijt} + \gamma_i \mathbb{1}(s_{it} = j) + \gamma^N \mathbb{1}(s_{it} = j) N_{it} + \phi^N N_{it} + \psi \cdot X_{ijt} + \epsilon_{ijt}. \quad (5)$$

The results are in Column (3) of Table 3. We do not find a significant coefficient for the interaction term, and thus conclude that the variety seeking we identified is not from new product discovery or investing in a promising option for the future. The robustness checks on learning and new product discovery correspond to the explore vs exploit tradeoff in models with information frictions. Appendix E.3 provides more details about the difference between variety-seeking behavior and the explore vs exploit tradeoff in a multi-armed bandit context.[21]

4.5.3 Rank Endogeneity

We observe the app version history of the platform and the factors included in the ranking algorithm. The platform provides to consumers a default comprehensive ranking that is based on four factors: price, distance, monthly sales, and ratings. We do not include distance and ratings in the baseline utility specification. A potential endogeneity problem is that if time-varying geographic distance and ratings enter both utility and ranking, then it will bias our position effect estimates away from zero. To investigate this possibility, we also estimate a model with distance and ratings entering the utility. Table 3 Column (4) shows the results. The distance coefficient is not significant at the 5% level, and the rating coefficient is not significant. The rank effect estimate does not change much. The insignificance of the rating coefficient is from a lack of variation both across kitchens and over time. The insignificance of the distance coefficient suggests that within the delivery distance limit, geographic distance

[20] An alternative definition of a new kitchen is a kitchen from which the consumer has never ordered. However, since there are kitchens that consumers do not like and will never order from, this definition would mix new kitchens and old disliked kitchens together.

[21] Other sources of information frictions include searching. However, search costs belong to forces that would motivate consumers to stay with the choices with which they are familiar with and from which they order repeatedly, which is inconsistent with our observation of frequent switching.

Table 4: Ranking Algorithms

	Ranking Algorithm
Random Ranking	Rank kitchens in a random order.
Platform Ranking	The platform ranking observed in the data.
Suboptimal Ranking	Rank kitchens in decreasing order of expected utility, without considering the variety seeking effect ($\gamma_i = 0$).
Optimal Ranking	Rank kitchens in decreasing order of expected utility, considering the variety seeking effect ($\gamma_i \neq 0$).

does not affect consumers' utility substantively. Thus, variation over time in rank associated with geographic location changes is exogenous.

5 Counterfactual Analysis of State-based Targeting

In this section, we use our demand estimates to study the managerial implications of variety seeking for targeted ranking and targeted pricing.

5.1 State-based Targeted Ranking

In this section, we examine the implications of variety-seeking preferences for targeted ranking. We use our demand estimates to compare consumer welfare, platform revenue, and purchase probability under several different ranking algorithms shown in Table 4.

For each ranking algorithm, we hold the choice set, prices, and all variables in X_{ijt} fixed as we observed in the data, and only change the ranking positions of kitchens. The random ranking benchmark ranks kitchens randomly.[22] The platform ranking uses the default ranking we observed in the real dataset. The suboptimal ranking ranks kitchens in decreasing order of expected utility, where the expected utility is calculated without the variety-seeking term. The optimal ranking ranks kitchens in the decreasing order of expected utility with the variety-seeking term. The optimal ranking optimizes ex-ante consumer welfare.[23]

We study the ranking effects allowing consumers to have heterogeneous state-dependent preferences, and the platform presents personalized rankings accounting for them. We mea-

[22]Specifically, the results are averaged across multiple randomly drawn rankings.

[23]Appendix D.1 provides a detailed discussion and technical proof of the optimality of rankings. Note that the optimal ranking of consumer welfare used in the analysis does not necessarily optimize platform revenue, for which there is no simple index like utility to follow, and one need to exhaust all the $J!$ potential combinatorics to compare revenues.

sure the total effect of ranking by the difference between the optimal ranking and the random ranking benchmark. The difference between the optimal ranking and the suboptimal ranking represents the effect of considering consumer variety-seeking preferences in ranking. Then the ratio of the two shows how much of the improvement from random ranking to optimal ranking is due specifically to the inclusion of variety-seeking preferences. We study this by taking 1000 draws from the estimated distribution of state-dependent preferences. For each draw of random coefficients, we simulate consumer choices 50 times to get the market outcome. For each market outcome, we calculate the ratio $\Delta S = \frac{S_{optimal} - S_{suboptimal}}{S_{optimal} - S_{random}}, S \in \{W, \Pi, q\}$ to measure the state-based targeted ranking effect on consumer welfare, platform revenue, and purchase probability. For each consumer i, $\Delta\tilde{S}$ is a function of the state-dependent preference γ_i, and the choice set effect C_i.[24] Equations 6, 7, and 8 present the average ranking effect on the platform.

$$\Delta W = \int \Delta \sum_i W(\gamma_i, C_i) dF(\gamma_i), \tag{6}$$

$$\Delta \Pi = \int \Delta \sum_i \Pi(\gamma_i, C_i) dF(\gamma_i), \tag{7}$$

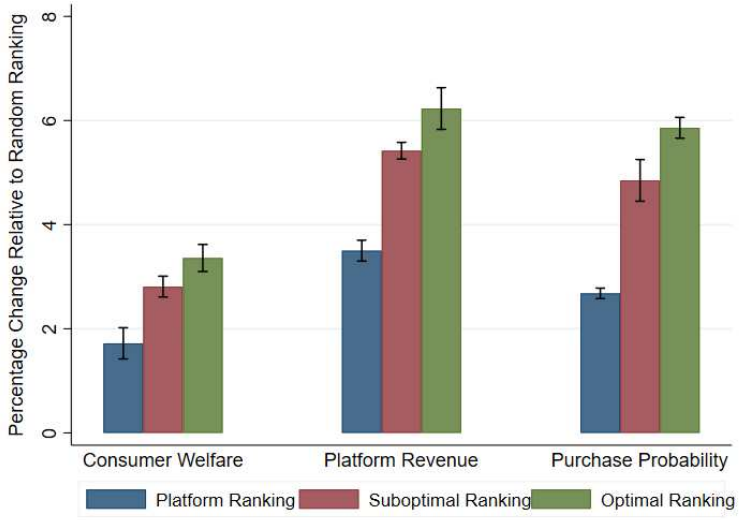
$$\Delta q = \int \frac{1}{N} \Delta \sum_i q(\gamma_i, C_i) dF(\gamma_i). \tag{8}$$

Figure 2 presents the average ranking effects on consumer welfare, platform revenue, and purchase probability. We find that optimizing the ranking algorithm with variety-seeking preferences takes up 14.6% of the revenue improvement, 17.2% of the consumer welfare improvement, and 18.3% of the purchase probability improvement, out of the total ranking effect. Figure 3 presents the detailed distribution of state-based targeted ranking effect on the platform level across simulations.

We further decompose the effect of optimal ranking on different market outcomes into higher match utility and lower position cost. Optimal ranking only generates different market outcomes from suboptimal ranking when the variety-seeking effects are high enough to either change the position of options or change the choice of consumers. Table 5 summarizes different scenarios when considering variety-seeking effects in ranking algorithms that generate different market outcomes, and shows how consumer welfare (through position cost and utility), platform revenue, and purchase probability are affected in each scenario. Specifically, we find that, on average, 8% of the increase in welfare comes from lower position costs, with the optimal ranking

[24] To understand the intuition why ΔS is a function of individual choice set C_i, consider an extreme example where there is no overlapping options of a consumer's choice set across different periods, then we the state-based targeted ranking will have no effect on this consumer's choices.

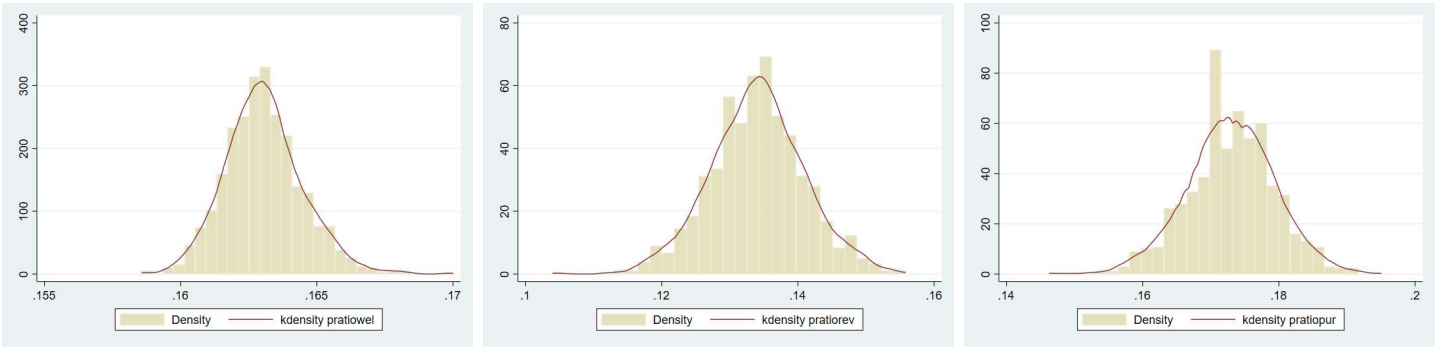
Figure 2: Ranking Effects with Heterogeneous State-Based Targeting



Note: The figure shows the percentage change in average platform-level welfare, revenue, and purchase probability of platform ranking, suboptimal ranking, and optimal ranking relative to the random ranking benchmark. We simulate choices under each given ranking algorithm 50 times to get the mean and standard deviation of the corresponding variables.

Figure 3: Distribution of Ranking Effects with Heterogeneous State-Based Targeting

(A) Welfare (B) Revenue (C) Purchase Probability



Note: The figure shows the detailed distribution of state-based targeted ranking effect on the platform level. Panel(A)-(C) presents the distribution of platform-level consumer welfare, revenue, and purchase probability across 1000 simulations.

decreasing the average position purchased by 0.2 positions.[25] Most of the welfare gain (92%) is from consumers purchasing a more desirable seller under the optimal ranking, which was ranked lower under the suboptimal ranking, such that the consumers choose the outside option or a

[25]Note that in the comparison between optimal and suboptimal ranking, the average position where a transaction happens is moved upwards by at most 1 position, since the utility index will be different for at most one option (the previous period choice) for each choice set.

Table 5: Decomposition of Optimal Ranking Effect

Choice under Supoptimal Ranking	Choice under Optimal Ranking	Market Outcome			
		Consumer Welfare		Revenue	Purchase Prob.
		Utility	Position Cost		
Outside Option	Different seller	↑	↑	↑	↑
Last period choice	Different seller	↑	–	Uncertain	–
Different seller	Different seller	–	↓	Uncertain	–

Note: The table shows the different scenarios where considering variety-seeking effects in ranking algorithms could affect consumer welfare (through position cost and utility), platform revenue, and purchase probability.

worse match ranked at a higher position. We also find that the two utility-based rankings are both better than the platform ranking.[26]

5.2 State-based Targeted Pricing

In this section, we use the demand estimates to investigate the implications of consumers' variety-seeking preferences for competitive targeted pricing. Given the large number of consumers and sellers in our data, it is intractable to estimate a full dynamic pricing game model. Instead, we use the data to calibrate parameters in consumer preference and simulate the optimal pricing strategy in a duopoly model, in order to understand the pricing implications under variety seeking. Consumers always choose with variety-seeking preferences, and we check the profit and welfare difference between optimal targeted pricing, where kitchens' policy functions are gained from the real state dependence parameter, and suboptimal pricing, where the kitchens do not consider the state dependence preference when making pricing decisions. We also analyze the comparative statics when the level of variety seeking varies.

The model consists of two kitchens competing for consumers with variety-seeking preferences by pricing heterogeneous products. Each kitchen sets a pricing policy to maximize the discounted sum of profits over an infinite horizon. Demand is derived from a population of consumers who make discrete choices from two products and an outside option. For simplicity, we drop the consumer-specific index. In each period t, a consumer has chosen one option in the last period, $s_t \in \mathcal{J} = \{0, 1, 2\}$. If the consumer chose option j in the last period, $s_t = j$, and

[26] In our sample, this could result from the naivety of the ranking algorithm the start-up platform is using in the early stage, which is not personalized except for accounting for the customers' distances to kitchens. Moreover, existing literature consistently documents that utility-based rankings increase consumer welfare and purchase probability.(Compiani et al., 2021; Ursu, 2018; De los Santos and Koulayev, 2017; Chen and Yao, 2017; Ghose et al., 2014, 2012). Results on how a utility-based ranking affects revenues are mixed. Ghose et al. (2014) find that it yields the largest total revenue among the rankings considered, whereas Ursu (2018) find that revenues decrease in three out of the four destinations in her sample.

purchases product, $k \in \mathcal{J}$, then his or her state becomes, $s_{t+1} = k$. Conditional on price p_{jt} and the consumer's current state s_t, the utility from the product j at time t is as follows:

$$u_{jt} = \beta_j + \alpha p_{jt} + \gamma 1\{s_t = j\} + \epsilon_{jt}. \tag{9}$$

The preference parameters are from our demand estimation. We assume that the random utility component, ϵ_{jt}, is i.i.d. Type I extreme value distributed. We study heterogeneous products and use the mean estimates μ_j of the estimates of β_j of the two most popular kitchen groups. Price is targeted based on the state:

$$p_{jt} = \begin{cases} p_{jt}^{own}, \text{ if } s_t = j. \\ p_{jt}^{rival}, \text{ if } s_t \neq j. \end{cases} \tag{10}$$

At any point in time, the market is summarized by the distribution of consumers over the states. Let $x_{jt} \in [0, 1]$ be the fraction of consumers who chose j in the last period. Let $U_j(s_t, p_t)$ denote the deterministic component of the utility index, such that $u_{jt} = U_j(s_t, p_t) + \epsilon_{jt}$. The consumer's choice probability has the following logit form:

$$P_j(s_t, \mathbf{p_t}) = \frac{exp[U_j(s_t, p_{jt})]}{\sum_{k \in \mathcal{J}} exp[U_k(s_t, p_{kt})]}. \tag{11}$$

We assume that a seller is able to observe whether a consumer ordered from its kitchen in the last period or not, and potentially sets different prices for its own consumers ($s_t = j$) and its rival seller's consumers ($s_t \neq j$). We obtain the respective aggregate demand by summing the demand over consumer states. Demand from own consumers and rival's consumers for product j are:

$$D_j^{own}(\mathbf{x_t}, \mathbf{p_t}) = x_{jt} P_j(j, \mathbf{p_t}) \tag{12}$$

$$D_j^{rival}(\mathbf{x_t}, \mathbf{p_t}) = \sum_{k \in \mathcal{J}/\{j\}} x_{kt} P_j(k, \mathbf{p_t}) \tag{13}$$

The distribution of consumer states, x_t, summarizes all current-period payoff-relevant information for the kitchen and describes the state of the market. The transition of the aggregate state can be derived from the transition probabilities of the individual states. Conditional on a price vector $\mathbf{p_t}$, we can define a Markov transition matrix $\mathbf{Q}(\mathbf{p_t})$ with the following elements:

$$Q_{jk}(\mathbf{p_t}) = P_j(k, \mathbf{p_t}). \tag{14}$$

The whole state vector evolves according to the Markov chain:

$$\mathbf{x_{t+1}} = \mathbf{Q}(\mathbf{p_t})\mathbf{x_t}. \tag{15}$$

The evolution of the state vector is deterministic, and we denote the transition function by f, $\mathbf{x_{t+1}} = f(\mathbf{x_t}, \mathbf{p_t})$. Time is discrete, $t = 0, 1, ...$ Conditional on all product prices and the state of the market, x_t, kitchen j's current period profit function is

$$\pi_j(\mathbf{x_t}, \mathbf{p_t}) = D_j^{own}(\mathbf{x_t}, \mathbf{p_t}) \cdot (p_j^{own} - c_j) + D_j^{rival}(\mathbf{x_t}, \mathbf{p_t}) \cdot (p_j^{rival} - c_j) \tag{16}$$

where c_j is the marginal cost of production, which does not vary over time. We use historical ingredient prices[27] during December 10 to 20, 2016, to calibrate the marginal cost of production, and adjust the variance of the logit distribution so that the equilibrium price at the estimated variety-seeking level equals the observed mean price in our data. Kitchens compete in prices and choose Markovian strategies, $\sigma_j : \mathbf{X} \to R^2$, that depend on the current payoff-relevant information, summarized by $\mathbf{x}$. This assumption rules out strategies that depend on historical prices and potential collusion among sellers. Kitchens discount the future using the common factor β, $0 < \beta < 1$. For a given profile of strategies, $\sigma = (\sigma_1, \sigma_2)$, the present discounted value of profits is $\sum_{t=0}^{\infty} \beta^t \pi_j[x_t, \sigma(x_t)]$. Conditional on a profile of competitors' strategies, σ_{-j}, kitchen j chooses a pricing strategy that maximizes its expected value. Associated with a solution to this problem is kitchen j's value function, which satisfies the Bellman equation

$$V_j(\mathbf{x}) = \max_{\mathbf{p_j} \geq 0} \{\pi_j(\mathbf{x}, \mathbf{p}) + \beta V_j[f(\mathbf{x}, \mathbf{p})]\}, \forall \mathbf{x} \in \mathbf{X}. \tag{17}$$

We use Markov Perfect Equilibrium (MPE) as our solution concept. In pure strategies, MPE is defined by a pricing strategy for each kitchen, σ_j^*, and an associated value function, V_j, such that

$$V_j(\mathbf{x}) = \max_{\mathbf{p_j} \geq 0} \{\pi_j[\mathbf{x}, \mathbf{p_j}, \sigma_{-\mathbf{j}}^*(\mathbf{x})] + \beta V_j\{f[\mathbf{x}, \mathbf{p_j}, \sigma_{-\mathbf{j}}^*(\mathbf{x})]\}\}, \tag{18}$$

for all states, $\mathbf{x}$, and kitchens. In each subgame starting at $\mathbf{x}$, the kitchen's strategy is a best response to the strategies its competitor chooses.[28]

We calculate steady-state prices as follows. We begin by computing the equilibrium pricing strategies of each kitchen. Then, we choose an arbitrary initial state for the period $t = 0$ and

[27] Historic prices of ingredients come from the National Bureau of Statistics website: `http://www.stats.gov.cn/tjsj/zxfb/201612/t20161226_1445711.html`

[28] Appendix B1 of Dubé et al. (2009) proves the existence of a pure strategy equilibrium of a similar model. Appendix D.2 discusses the procedure for the numerical solution to the dynamic programming problem.

Table 6: Uniform Pricing v.s. Targeted Pricing

Demand	Supply	
Consumers choose with VS preferences($\gamma = \hat{\gamma} < 0$)	(1)Uniform Pricing (Status quo)	Sellers set the same price to all consumers ($\gamma = 0, p_j^{own} = p_j^{rival}$)
	(2)Targeted Pricing	Sellers set targeted prices based on last period choice ($\gamma = \hat{\gamma} < 0, p_j^{own} \neq p_j^{rival}$)

Note: The table presents details of how we compare the performance of two pricing schemes in the price competition equilibrium: uniform pricing and targeted pricing.

calculate the corresponding sequence of equilibrium price levels and state vectors for periods $t = 1, 2, ...$ We stop this process after the convergence of the state vector and corresponding equilibrium prices to fixed values occurs.

5.2.1 Results

Optimal Targeted Pricing We study the effect of a state-based targeted pricing strategy by comparing the performance of two pricing schemes: uniform pricing and targeted pricing. Table 6 presents the details of each pricing strategy. In both pricing strategies, consumers always choose with variety-seeking preferences. Uniform pricing corresponds to the case where kitchens do not consider consumers' variety-seeking preferences and do not do targeted pricing based on them. Optimal targeted pricing corresponds to the case when kitchens consider consumers' variety-seeking preferences and set different prices to their own and rival's consumers. Table 7 shows the percentage changes of the equilibrium steady state prices, profits, and consumer welfare. With optimal targeted pricing, kitchens set lower prices for their own consumers and higher prices for the rival seller's consumers; this strategy is similar to a quantity discount for existing customers. Optimal pricing increases average transaction prices by 6.8%, profits by 5.2%, and decreases consumer welfare by 9.1%.

This optimal pricing counterfactual predicts that it is profitable to do targeted pricing based on consumers' variety-seeking preferences. However, the platform we study and the most popular food delivery platforms do not have targeted pricing based on recent orders. The common existing industry practices in targeted pricing are focused on new consumer acquisition, past consumer retention, and whole platform promotion. Appendix E.1 provides more details about the current industry practice for targeted pricing.

Comparative Statics The effect of variety-seeking preferences on price competition is theoretically ambiguous, and empirical work is necessary to pin down the prediction. Besides

Table 7: Optimal Pricing

	% Change with Targeted Pricing
Price_{1own}	-0.72
Price_{1rival}	9.45
Price_{2own}	-0.20
Price_{2rival}	9.54
Average Price	6.91
Total Profits	5.23
Platform Revenue	1.75
Consumer Welfare	-9.22

Note: Percentage change is calculated by formula $\frac{Targeted-Uniform}{Uniform}$ of the corresponding variable. The average price is the weighted sum of prices, where weights are the corresponding choice probabilities.

the comparison between optimal targeted pricing and uniform pricing, we also study how the optimal targeted pricing levels change with the variety-seeking level. We compute equilibrium prices for a range of variety seeking achieved by scaling the mean value of γ_i gained from demand estimation $s \in \{0, \frac{1}{3}, \frac{1}{2}, 1, 2, 3\}$.[29] Figure 4 illustrates the key pricing incentives in our model. The x-axis is the level of variety-seeking preferences represented by a scale factor. Similar to the traditional harvest vs invest tradeoff in a switching cost model, there is also a current period vs future incentive difference under variety seeking. Panels (A) and (C) in Figure 4 show kitchen 1's prices for its own and rival's consumers, respectively. Panels (B) and (D) show kitchen 2's prices for its own and rival's consumers. To decompose kitchens' current and future period incentives, in each graph we plot the equilibrium price levels for both cases when kitchens are forward-looking ($\beta = 0.98$) and myopic ($\beta = 0$). Then the difference between the price levels when $\gamma = 0$ and the myopic line is the current period incentive for kitchens, and the difference between the price levels when kitchens are myopic and those when kitchens are forward-looking reflects the future incentive of kitchens. In terms of the current period incentive, for consumers who purchased from the kitchen in the last period, the kitchen has the incentive to decrease the price to keep them from switching to the competitor. For consumers who did not purchase from the kitchen in the last period, the kitchen has the incentive to increase the price to exploit the market power induced by consumers' variety-seeking preferences. In terms of future incentives, the tendency to switch from consumers' variety-seeking preferences discourages kitchens from competing for them and imposes upward pressure on prices. To understand this more formally,

[29]The corresponding level of variety seeking is the mean estimate of γ_i multiplied by the scale factor $\gamma = \hat{\gamma} \cdot s$. The variety-seeking level we estimate from the original data corresponds to a scale factor equal to one. The equilibrium when consumers don't have state-dependent preferences corresponds to a scale factor of zero.

Table 8: Variety Seeking's Effect on Kitchens' Pricing Incentives

	Price for Own Consumer	Price for Rival's Consumer
Current Period	–	+
Future	+	+

Note: The table summarizes the effect of consumers' variety-seeking preferences on the kitchen's pricing incentives for its own and rival's consumers in the current period and future. '+' suggests an upward pressure on price levels, and '–' suggests a downward pressure on prices.

we can write down seller j's profit-maximizing first-order conditions based on equation 18:

$$\underbrace{\frac{\partial \pi_j}{\partial p_j^{own}}}_{\text{Current profits}} + \underbrace{\beta \frac{\partial V_j\{f[\mathbf{x}, \mathbf{p_j}, \sigma^*_{-\mathbf{j}}(\mathbf{x})]\}}{\partial p_j^{own}}}_{\text{Future distribution of consumer states}} = 0, \tag{19}$$

$$\underbrace{\frac{\partial \pi_j}{\partial p_j^{rival}}}_{\text{Current profits}} + \underbrace{\beta \frac{\partial V_j\{f[\mathbf{x}, \mathbf{p_j}, \sigma^*_{-\mathbf{j}}(\mathbf{x})]\}}{\partial p_j^{rival}}}_{\text{Future distribution of consumer states}} = 0. \tag{20}$$

The prices affect the current period profits, and meanwhile affect the future distribution of consumer states. Table 8 provides a summary of the direction of the effect of consumers' variety-seeking preferences on kitchens' pricing incentive. When kitchens are very differentiated, it is possible that variety seeking can intensify price competition, and the driving force comes from the more popular kitchens trying to maintain a large own consumer base.[30] But in our demand estimate range, it turns out that kitchens are not differentiated enough, and consumers' variety-seeking preferences soften kitchen price competition.

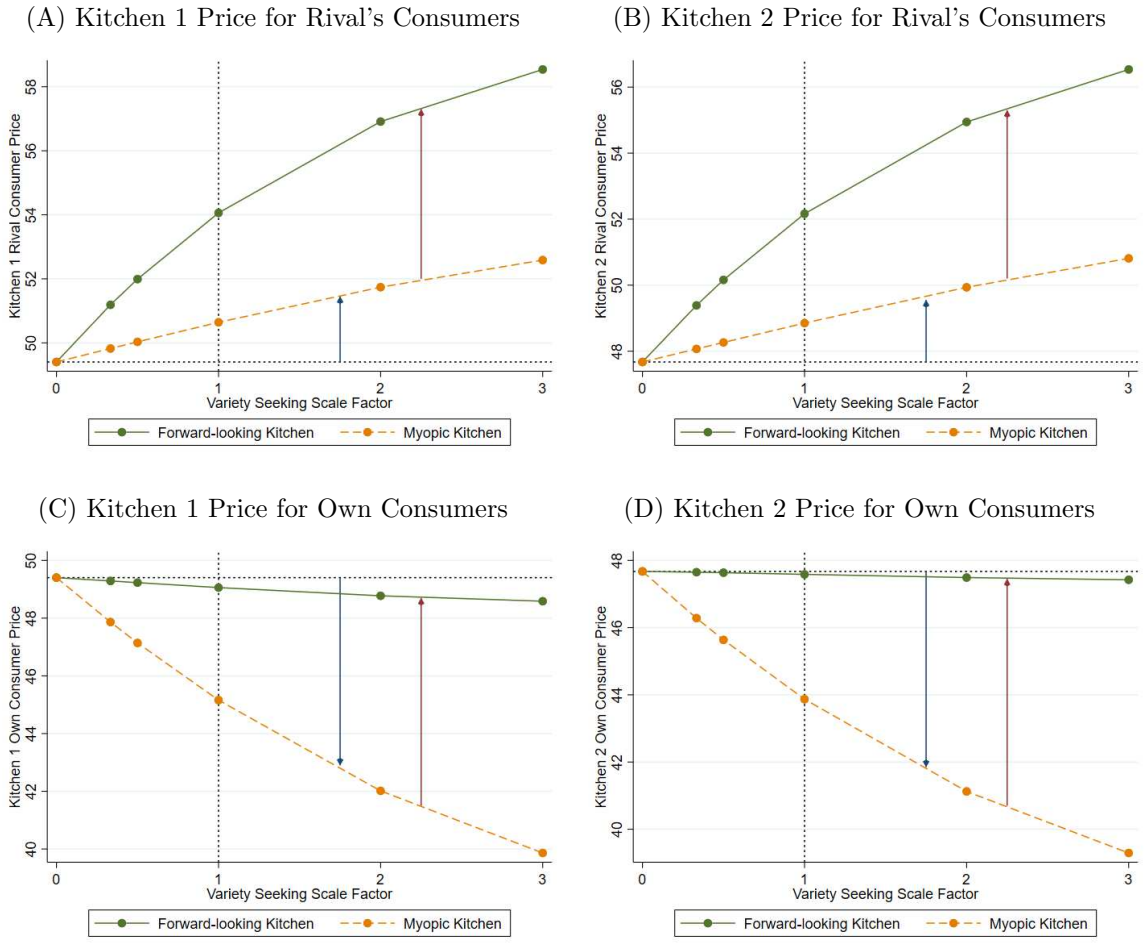
In order to study the effect of variety-seeking preferences on equilibrium profit and welfare level and eliminate the difference between "staying cost" and "switching bonus" models, which is due to the differential impact on the outside good market share under changes in γ, we calculate the difference of the steady state profit and welfare relative to a pure demand response benchmark.

Denote the steady state equilibrium price, market shares, profit, and welfare level when consumers' variety-seeking preferences equals γ as $\{\mathbf{p}_\gamma, \mathbf{x}_\gamma(\mathbf{p}_\gamma), \pi(\mathbf{p}_\gamma, \mathbf{x}_\gamma(\mathbf{p}_\gamma)), w(\mathbf{p}_\gamma, \mathbf{x}_\gamma(\mathbf{p}_\gamma), \gamma)\}$. The equilibrium profits are defined by the total profits of all firms

$$\pi(\mathbf{p}_\gamma, \mathbf{x}_\gamma(\mathbf{p}_\gamma)) = \sum_{j \in \mathcal{J}/\{0\}} \pi_j(\mathbf{p}_\gamma, \mathbf{x}_\gamma(\mathbf{p}_\gamma))$$

[30] Appendix D.3 provides more intuition for how product differentiation and variety-seeking level affect equilibrium price competition level.

Figure 4: Steady State Equilibrium Pricing



Note: The figure shows the steady state equilibrium prices for different levels of variety-seeking preferences. The x-axis is the scale factor, and the corresponding variety-seeking level equals the scale factor times the mean estimate of γ_i. Panels (A) and (C) show Kitchen 1's prices for its rival's and its own consumers, respectively. Panels (B) and (D) show Kitchen 2's prices for its rival's and its own consumers. In each graph, we plot the equilibrium price levels for cases when kitchens are forward-looking ($\beta = 0.98$) and myopic ($\beta = 0$). The difference between the price level at $\gamma = 0$ and the myopic kitchen price level reflects the kitchens' current period pricing incentives (blue arrow), and the difference between myopic and forward-looking kitchen prices captures the kitchens' future pricing incentives (red arrow).

The equilibrium welfare is defined by the sum of ex-ante welfare of consumers across all states.

$$w(\mathbf{p}_\gamma, \mathbf{x}_\gamma(\mathbf{p}_\gamma), \gamma) = -\frac{1}{\alpha} \sum_{j \in \mathcal{J}} x_j \cdot \left[ln \left(\sum_{k \in \mathcal{J}} exp(U_k(j, p_k)) \right) \right]$$

where x_j denotes the share of consumers with $s_t = j$, and p_k denotes product k's price for consumers with $s_t = j$. Table 9 presents the decomposition of welfare and profit effects of changing the levels of consumers' variety-seeking preference. Starting from the steady state equilibrium when consumers do not have state-dependent preference ($\{\pi(\mathbf{p}_0, \mathbf{x}_0(\mathbf{p}_0)), w(\mathbf{p}_0, \mathbf{x}_0(\mathbf{p}_0), 0)\}$), as γ

Table 9: Decomposition of Welfare and Profit Effects of Variety-Seeking Preferences

Decomposition	Market Outcome
(0) No VS	$\{\pi(\mathbf{p}_0, \mathbf{x}_0(\mathbf{p}_0)), w(\mathbf{p}_0, \mathbf{x}_0(\mathbf{p}_0), 0)\}$
(1) VS, demand-preference	$\{\pi(\mathbf{p}_0, \mathbf{x}_0(\mathbf{p}_0)), w(\mathbf{p}_\gamma, \mathbf{x}_0(\mathbf{p}_0), \hat{\gamma})\}$
(2) VS, demand-choice	$\{\pi(\mathbf{p}_0, \mathbf{x}_{\hat{\gamma}}(\mathbf{p}_0)), w(\mathbf{p}_0, \mathbf{x}_{\hat{\gamma}}(\mathbf{p}_0), \hat{\gamma})\}$
(3) VS, supply and equilibrium	$\{\pi(\mathbf{p}_{\hat{\gamma}}, \mathbf{x}_{\hat{\gamma}}(\mathbf{p}_{\hat{\gamma}})), w(\mathbf{p}_{\hat{\gamma}}, \mathbf{x}_{\hat{\gamma}}(\mathbf{p}_{\hat{\gamma}}), \hat{\gamma})\}$

Note: The table shows decomposition of equilibrium analysis when consumers have varying levels of variety-seeking preferences. Counterfactual (0) corresponds to the steady state equilibrium when consumers don't have state-dependent preferences. Starting from this equilibrium, counterfactual (1), (2), and (3) denote the market outcome when (1) only consumers change state-dependent preferences to $\gamma = \hat{\gamma}$, (2) consumers changes preferences and choices when sellers don't change prices, (3) sellers change prices and consumers change choices responding to the new price levels.

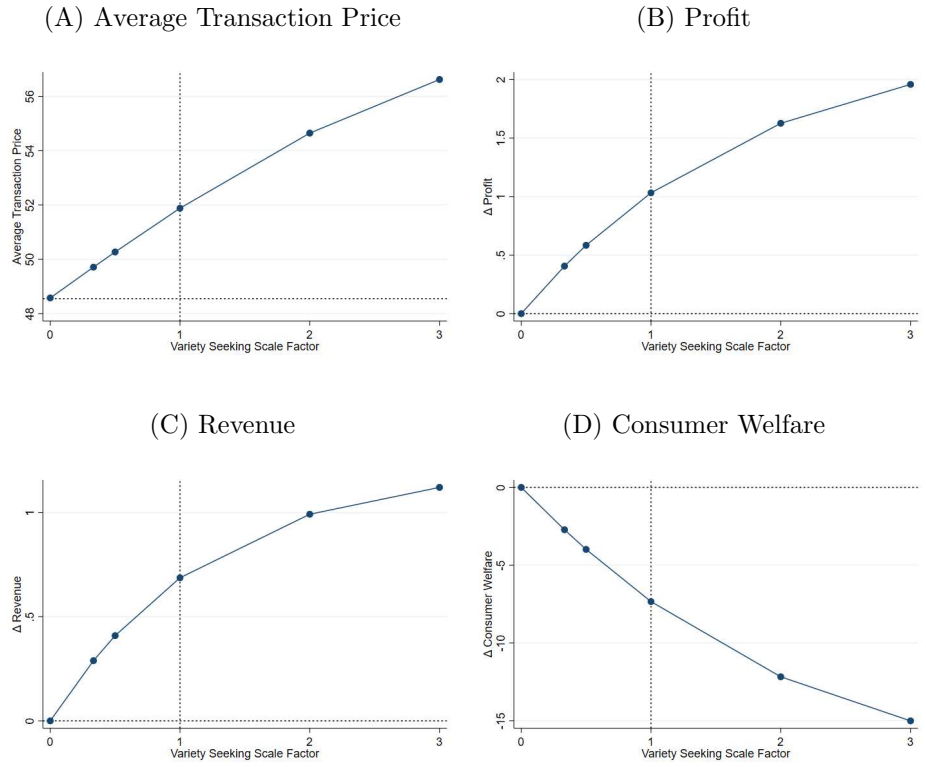
becomes negative, we decompose the changes in the steady state equilibrium into three channels. First, before prices and choices change, consumer welfare decreases because of the "staying cost" imposed by boredom. Second, before prices change, as γ decreases the utility from consumption, consumers will switch to the outside option, which further decreases profits and increases consumer welfare. Third, kitchens change price levels given the new consumer state-dependent preferences, and consumers respond to the changes. The channel of changes in equilibrium that we are interested in is the third channel, which captures the effect of consumers' variety-seeking preferences on the price competition level. The first two channels are pure demand side responses from changes in γ and are susceptible to the model specification of "staying cost" and "switching bonus".[31] This challenge of welfare and profit analysis when consumer preferences change is related to the challenge in welfare analysis of persuasive advertising discussed in Becker and Murphy (1993).

To tackle this problem, we use the pure demand response benchmark to separate the price competition effects on profits and consumer welfare from the pure demand response. The profit and welfare in the pure demand response benchmark for each variety-seeking level are calculated by equations 12 and 13, where prices and initial state variables are at the $\gamma = 0$ steady state equilibrium level, i.e., the pure demand response benchmark profit and welfare are $\pi(p_0, \gamma)$ and $u(p_0, \gamma)$, respectively. Then the price competition effect at each variety-seeking level is $\Delta\pi = \pi(p_\gamma, \gamma) - \pi(p_0, \gamma), \Delta u = u(p_\gamma, \gamma) - u(p_0, \gamma)$.[32]

[31] In this paper, we used the "staying cost" specification in accordance with boredom and satiation. The "switching bonus" specification would be more consistent with thrill seeking. We do not try to identify the different psychological factors behind negative state-dependent preference in this project.

[32] Dubé et al. (2009) adopts a different approach to solve this problem by changing the intercept of the outside option in order to keep the outside option market share fixed. This method also removes the extensive margin demand response to price changes in price competition in the third channel. We adopt an alternative method

Figure 5: Steady State Equilibrium



Note: The figure shows steady state equilibrium average transaction price, profit, revenue, and consumer welfare. The average transaction price is the weighted average of prices by transaction probability. Profits, revenue, and welfare are measured in CNY and are calculated relative to the pure demand response benchmark to measure the pure price competition effect.

Figure 5 shows how steady state prices, profits, and consumer welfare change with the level of consumers' variety-seeking preferences. We find that as the variety-seeking level increases, equilibrium average transaction prices and profits increase, whereas consumer welfare decreases. To summarize, we find that consumers' variety-seeking preferences soften price competition.

6 Conclusion

In this paper, we document consumers' variety-seeking preferences in high-frequency consumption. We show that a substantial fraction of consumers have negative state-dependent preferences after controlling for potential confounding factors. We also further ruled out alternative sources of structural state dependence, including learning and new product discovery. We find that consumers, on average, are willing to pay 23.1% more for switching to a different seller.

here to capture the extensive margin effect of price competition level.

There is heterogeneity in consumers' state dependence preference; 29.5% of consumers have inertia, whereas 70.5% of consumers exhibit variety-seeking preferences. This finding has important managerial implications for targeted marketing.

Using the demand estimates, we examine the managerial implications of variety seeking on targeted ranking and targeted pricing. For targeted ranking, we examine the profits, welfare, and purchase probability of several different ranking schemes. We find that optimizing the ranking algorithm with variety-seeking preferences takes up 14.6% of the revenue improvement, 17.2% of the consumer welfare improvement, and 18.3% of the purchase probability improvement, out of the total ranking effect. For targeted pricing, we study the effect of optimal pricing with variety seeking and the comparative statics of the effect of consumers' variety-seeking preferences on price competition. We find that optimal targeted pricing increases prices and profits, and decreases consumer welfare. The comparative statics analysis shows that the average equilibrium prices and profits increase with consumers' variety-seeking level, whereas consumer welfare decreases with it. In general, variety seeking softens price competition.

Although we only do counterfactual analysis on pricing and ranking in this paper, consumers' variety-seeking preferences have more general managerial implications in other areas, including targeted advertising, personalized recommendation systems (Ansari et al., 2000; Ansari and Mela, 2003), positioning (Sajeesh and Raju, 2010; Zeithammer and Thomadsen, 2013), and the distributional channel.[33] The managerial implications of consumer variety-seeking preferences are also not limited to the food delivery industry. Cadario and Morewedge (2021) documents the phenomenon that many people eat the same breakfast every day, yet seek variety for lunch and dinner, and attribute this difference to a psychological driver: variance in the pursuit of hedonic and utilitarian goals across meals. This is also helpful to understand the direction of state dependence we might expect in different industries. In utilitarian choice situations, such as workday breakfast, health care, gasoline, and grocery products, previous research has found evidence of inertia. In hedonic choice situations, such as lunch, dinner, music, movies, podcasts, magazines, art, and resorts, where people care more about entertainment and pleasure, it is important to consider consumers' variety-seeking preferences.

References

Abaluck, J., G. Compiani, and F. Zhang (2025). A method to estimate discrete choice models that is robust to consumer search. *Journal of Political Economy*, forthcoming.

[33] Cosguner et al. (2017) and Cosguner et al. (2018) study dynamic pricing models with switching cost in a distribution channel. The implications of variety seeking in a vertical relationship are also important to investigate.

Ackerberg, D. A. (2003). Advertising, learning, and consumer choice in experience good markets: an empirical examination. *International Economic Review 44*(3), 1007–1040.

Ansari, A., S. Essegaier, and R. Kohli (2000). Internet recommendation systems.

Ansari, A. and C. F. Mela (2003). E-customization. *Journal of marketing research 40*(2), 131–145.

Bass, F. M., M. M. Givon, M. U. Kalwani, D. Reibstein, and G. P. Wright (1984). An investigation into the order of the brand choice process. *Marketing Science 3*(4), 267–287.

Bawa, K. (1990). Modeling inertia and variety seeking tendencies in brand choice behavior. *Marketing Science 9*(3), 263–278.

Becker, G. S. and K. M. Murphy (1993). A simple theory of advertising as a good or bad. *The Quarterly Journal of Economics 108*(4), 941–964.

Besanko, D., J.-P. Dubé, and S. Gupta (2003). Competitive price discrimination strategies in a vertical channel using aggregate retail data. *Management Science 49*(9), 1121–1138.

Bonhomme, S., T. Lamadon, and E. Manresa (2019). A distributional framework for matched employer employee data. *Econometrica 87*(3), 699–739.

Bonhomme, S., T. Lamadon, and E. Manresa (2022). Discretizing unobserved heterogeneity. *Econometrica 90*(2), 625–643.

Broda, C. and D. E. Weinstein (2006). Globalization and the gains from variety. *The Quarterly journal of economics 121*(2), 541–585.

Brynjolfsson, E., Y. Hu, and M. D. Smith (2003). Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management science 49*(11), 1580–1596.

Cadario, R. and C. K. Morewedge (2021). Why do people eat the same breakfast every day? goals and circadian rhythms of variety seeking in meals. *Appetite*, 105716.

Che, H., S. Seetharaman, and K. Sudhir (2003). Pricing behavior in markets with state-dependence in demand. *Available at SSRN 442580*.

Chen, Y. and S. Yao (2017). Sequential search with refinement: Model and application with click-stream data. *Management Science 63*(12), 4345–4365.

Chintagunta, P. K. (1998). Inertia and variety seeking in a model of brand-purchase timing. *Marketing Science 17*(3), 253–270.

Chintagunta, P. K. (1999). Variety seeking, purchase timing, and the "lightning bolt" brand choice model. *Management Science 45*(4), 486–498.

Compiani, G., G. Lewis, S. Peng, and P. Wang (2024). Online search and optimal product rankings: An empirical framework. *Marketing Science 43*(3), 615–636.

Compiani, G., G. Lewis, S. Peng, and W. Wang (2021). Online search and product rankings: A double index approach. *Available at SSRN 3898134*.

Cosguner, K., T. Y. Chan, and P. Seetharaman (2017). Behavioral price discrimination in the presence of switching costs. *Marketing Science 36*(3), 426–435.

Cosguner, K., T. Y. Chan, and P. Seetharaman (2018). Dynamic pricing in a distribution channel in the presence of switching costs. *Management Science 64*(3), 1212–1229.

Datta, H., G. Knox, and B. J. Bronnenberg (2018). Changing their tune: How consumers' adoption of online streaming affects music consumption and discovery. *Marketing Science 37*(1), 5–21.

De los Santos, B. and S. Koulayev (2017). Optimizing click-through in online rankings with endogenous search refinement. *Marketing Science 36*(4), 542–564.

Derakhshan, M., N. Golrezaei, V. Manshadi, and V. Mirrokni (2022). Product ranking on online platforms. *Management Science*.

Donnelly, R., F. J. Ruiz, D. Blei, and S. Athey (2021). Counterfactual inference for consumer choice across many product categories. *Quantitative Marketing and Economics*, 1–39.

Dubé, J.-P., G. J. Hitsch, and P. E. Rossi (2009). Do switching costs make markets less competitive? *Journal of Marketing research 46*(4), 435–445.

Dubé, J.-P., G. J. Hitsch, and P. E. Rossi (2010). State dependence and alternative explanations for consumer inertia. *The RAND Journal of Economics 41*(3), 417–445.

Dubé, J.-P., G. J. Hitsch, P. E. Rossi, and M. A. Vitorino (2008). Category pricing with state-dependent utility. *Marketing Science 27*(3), 417–429.

Eggers, A. C., G. Tuñón, and A. Dafoe (2021). Placebo tests for causal inference. *Working paper*.

Ernst, M. D. (2004). Permutation methods: a basis for exact inference. *Statistical Science*, 676–685.

Ershov, D. (2020). Consumer product discovery costs, entry, quality and congestion in online markets. *Unpublished manuscript*.

Farrell, J. and P. Klemperer (2007). Coordination and lock-in: Competition with switching costs and network effects. *Handbook of industrial organization 3*, 1967–2072.

Ghose, A., P. G. Ipeirotis, and B. Li (2012). Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content. *Marketing Science 31*(3), 493–520.

Ghose, A., P. G. Ipeirotis, and B. Li (2014). Examining the impact of ranking on consumer behavior and search engine revenue. *Management Science 60*(7), 1632–1654.

Givon, M. (1984). Variety seeking through brand switching. *Marketing Science 3*(1), 1–22.

Good, P. (2013). *Permutation tests: a practical guide to resampling methods for testing hypotheses*. Springer Science & Business Media.

Hamilton, W., Z. Ying, and J. Leskovec (2017). Inductive representation learning on large graphs. *Advances in neural information processing systems 30*.

Heckman, J. J. et al. (1981). Heterogeneity and state dependence. *Studies in labor markets 31*, 91–140.

Holtz, D., B. Carterette, P. Chandar, Z. Nazari, H. Cramer, and S. Aral (2020). The engagement-diversity connection: Evidence from a field experiment on spotify. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pp. 75–76.

Howell, J. R., S. Lee, and G. M. Allenby (2016). Price promotions in choice models. *Marketing Science 35*(2), 319–334.

Jain, A., I. Liu, A. Sarda, and P. Molino (2019). Food discovery with uber eats: Using graph learning to power recommendations. *Uber engineering.*.

Jeziorski, P. and I. Segal (2015). What makes them click: Empirical analysis of consumer demand for search advertising. *American Economic Journal: Microeconomics 7*(3), 24–53.

Kahn, B. E. and J. S. Raju (1991). Effects of price promotions on variety-seeking and reinforcement behavior. *Marketing Science 10*(4), 316–337.

Kong, X., J.-P. Dubé, and O. Daljord (2022). Non-parametric estimation of habitual brand loyalty. *Available at SSRN 4070747*.

Koren, Y., R. Bell, and C. Volinsky (2009). Matrix factorization techniques for recommender systems. *Computer 42*(8), 30–37.

Larsen, B., C. H. Lu, and A. L. Zhang (2021). Intermediaries in bargaining: Evidence from business-to-business used-car inventory negotiations. Technical report, National Bureau of Economic Research.

Levine, J. and S. Seiler (2023). Identifying state dependence in brand choice: Evidence from hurricanes. *Marketing Science 42*(5), 934–957.

Luo, X., S. Tong, Z. Lin, and C. Zhang (2021). The impact of platform protection insurance on buyers and sellers in the sharing economy: A natural experiment. *Journal of Marketing 85*(2), 50–69.

MacKay, A. and M. Remer (2019). Consumer inertia and market power. *Available at SSRN 3380390*.

McAlister, L. (1982). A dynamic attribute satiation model of variety-seeking behavior. *Journal of Consumer Research 9*(2), 141–150.

McAlister, L. and E. Pessemier (1982). Variety seeking behavior: An interdisciplinary review. *Journal of Consumer research 9*(3), 311–322.

Miller, J. B. and A. Sanjurjo (2018). Surprised by the hot hand fallacy? a truth in the law of small numbers. *Econometrica 86*(6), 2019–2047.

Miller, J. B., A. Sanjurjo, et al. (2014). A cold shower for the hot hand fallacy. Technical report, Working paper.

Morozov, I., S. Seiler, X. Dong, and L. Hou (2021). Estimation of preference heterogeneity in markets with costly search. *Marketing Science 40*(5), 871–899.

Narayanan, S. and P. Manchanda (2009). Heterogeneous learning and the targeting of marketing communication for new products. *Marketing science 28*(3), 424–441.

Natan, O. R. (2020). Choice frictions in large assortments. Technical report.

Natan, O. R. (2025). Choice frictions in large assortments. *Marketing Science 44*(3), 593–625.

Osborne, M. (2011). Consumer learning, switching costs, and heterogeneity: A structural examination. *Quantitative Marketing and Economics 9*(1), 25–70.

Pancras, J. and K. Sudhir (2007). Optimal marketing strategies for a customer data intermediary. *Journal of Marketing research 44*(4), 560–578.

Park, S. and S. Gupta (2011). A regime-switching model of cyclical category buying. *Marketing Science 30*(3), 469–480.

Quan, T. W. and K. R. Williams (2018). Product variety, across-market demand heterogeneity, and the value of online retail. *The RAND Journal of Economics 49*(4), 877–913.

Rossi, P. E., R. E. McCulloch, and G. M. Allenby (1996). The value of purchase history data in target marketing. *Marketing Science 15*(4), 321–340.

Sajeesh, S. and J. S. Raju (2010). Positioning and pricing in a variety seeking market. *Management science 56*(6), 949–961.

Seetharaman, P. and H. Che (2009). Price competition in markets with consumer variety seeking. *Marketing Science 28*(3), 516–525.

Seetharaman, P. and P. Chintagunta (1998). A model of inertia and variety-seeking with marketing variables. *International Journal of Research in Marketing 15*(1), 1–17.

Simonov, A., J.-P. Dubé, G. Hitsch, and P. Rossi (2020). State-dependent demand estimation with initial conditions correction. *Journal of Marketing Research 57*(5), 789–809.

Skiti, T., X. Luo, and Z. Lin (2022). When more is less: Quality and variety trade-off in sharing economy platforms. *Journal of Management Studies 59*(7), 1817–1838.

Smith, A. N., Y. Li, S. Hasegawa, N. Terui, and G. M. Allenby (2020). Integrated factor models of variety seeking dynamics.

Smith, A. N., S. Seiler, and I. Aggarwal (2023). Optimal price targeting. *Marketing Science 42*(3), 476–499.

Tae, C. J., X. Luo, and Z. Lin (2020). Capacity-constrained entrepreneurs and their product portfolio size: The response to a platform design change on a chinese sharing economy platform. *Strategic Entrepreneurship Journal 14*(3), 302–328.

Trindade, A. and F. Caldieraro (2025). The effect of competition on retail price and assortment size decisions: Theory and evidence. *Journal of Retailing*.

Trivedi, M., F. M. Bass, and R. C. Rao (1994). A model of stochastic variety-seeking. *Marketing Science 13*(3), 274–297.

Ursu, R. M. (2018). The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science 37*(4), 530–552.

Van Herpen, E. and R. Pieters (2002). The variety of an assortment: An extension to the attribute-based approach. *Marketing Science 21*(3), 331–341.

Villas-Boas, J. M. (2015). A short survey on switching costs and dynamic competition. *International Journal of Research in Marketing 32*(2), 219–222.

Wang, Y., Y. Ning, I. Liu, and X. X. Zhang (2018). Food discovery with uber eats: recommending for the marketplace.

Yoganarasimhan, H. (2020). Search personalization using machine learning. *Management Science 66*(3), 1045–1070.

Zeithammer, R. and R. Thomadsen (2013). Vertical differentiation with variety-seeking consumers. *Management Science 59*(2), 390–401.

Zhang, J. and L. Krishnamurthi (2004). Customizing promotions in online stores. *Marketing science 23*(4), 561–578.

Zhang, J. and M. Wedel (2009). The effectiveness of customized promotions in online and offline stores. *Journal of marketing research 46*(2), 190–206.

A Platform Background Details

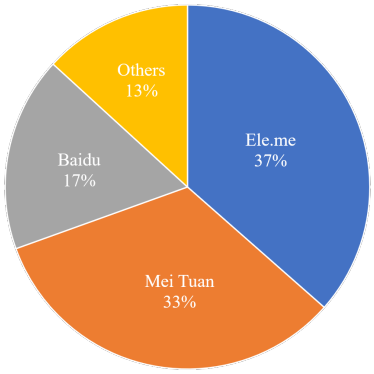
Market Structure Figure A.1 (A) shows the market share of the food delivery market in China in 2017. The total market value is $28 billion.[34] The platform we study has a relatively small market share and is a niche platform. Homemade food in China is believed to be cleaner and healthier than that in restaurants, given the food security scandals. Figure A.1 (B) shows the buyer and seller distribution map in the city of Beijing. The orange dots are the sellers,

[34] https://www.iimedia.cn/c400/60449.html

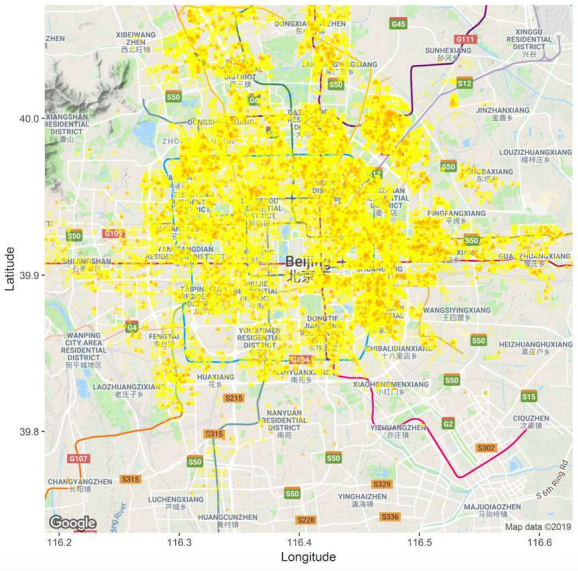
and the yellow dots are the buyers. In terms of delivery service, sellers can choose to deliver by themselves or use full-time third-party carriers. The delivery fees are equally shared by kitchens and consumers. [35]

Figure A.1: Food Delivery Market

(A) Market Structure of the Food Delivery Market



(B) Buyer and Seller Geographic Distribution



Note: Panel (A) shows the market share of the food delivery market in China, 2017. The platform we are studying is a relatively niche platform. Panel (B) shows the geographic distribution of sellers and buyers on the map. The orange dots are the sellers, and the yellow dots are the buyers.

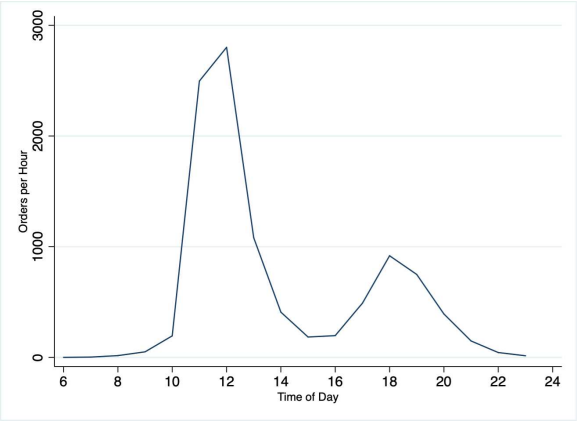
Daily Order Trend Figure A.2 shows the order trend during the day. Most orders are workday lunch orders. So we assume one consumption per period. We also observe the utensils in each order, and in 90% of our orders requested one set of utensils. This is also an advantage of our data since one major concern of using household-level packaged good scanner data is the confound of intrahousehold heterogeneity: multiple family members could be using the same membership card to purchase grocery products.

App Display The homepage of the platform for the consumer side is shown in Figure A.3. On the homepage, there is a kitchen list showing information about monthly sales, ratings, delivery radius, distance, address, seller's hometown (which is correlated with the kitchen's cuisine type), and pictures of the dishes, where the users can swipe left and right to browse the

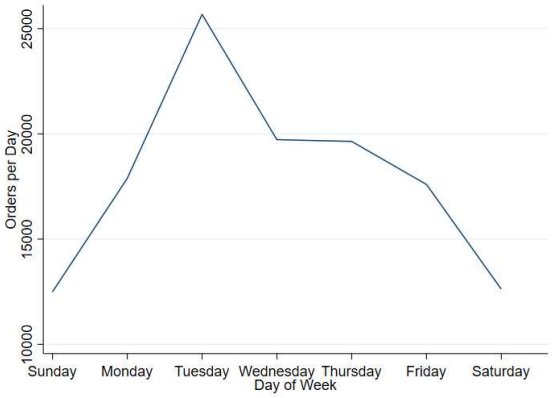
[35] More information about the platform background can be found in Tae et al. (2020), Luo et al. (2021), and Skiti et al. (2022).

Figure A.2: Order Trend

(A) Order Trend during the Day



(B) Order Trend during the Week



Note: The figure shows the timeline of orders during the day and the week. Panel (A) shows the number of orders per hour throughout the day. Panel (B) shows the number of orders per day on each day of the week. 80 % of orders are workday orders.

menu without clicking on the kitchen. After clicking on the detailed kitchen page, consumers can find out more about the menu and specific dishes, including prices, sales volume, and top reviews. More information about the kitchen can be found on the detailed kitchen information page, including the seller's demographic information, like hometown and age, time to join the platform, kitchen address, certificate, and kitchen reviews.

Cuisine Distribution We divide all kitchens into eleven different cuisines, where the distribution of each cuisine is in Figure A.4. The cuisine variable is constructed based on kitchen features, dish features, and hometown information of the kitchen owners.

HHI Distribution We compute the Herfindahl-Hirschman Index (HHI) as a measure of the dispersion level of consumers' chosen sets aggregated over time. Each consumer is seen as a market, and we define the index as $HHI_i = \sum_j ms_{ij}^2 \cdot 10000$, where ms_{ij} is the market share of orders from kitchen j among all consumer i's orders on the platform over time. The distribution of HHI_i in our data is shown in panel (B) of Figure A.4. The figure shows that most consumers' chosen sets are rather dispersed.

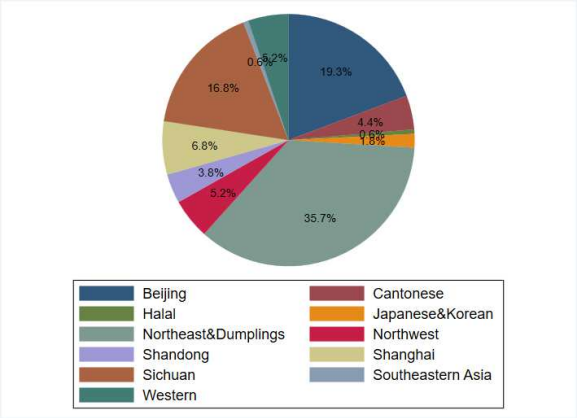
Figure A.3: Homepage Display



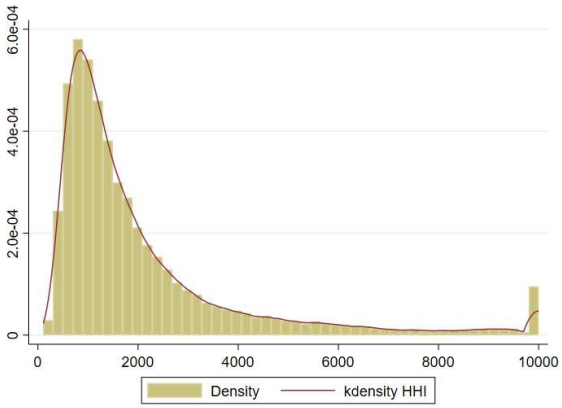
Note: The figure shows the platform homepage display. After clicking on a kitchen from the kitchen list, consumers can see more detailed dish lists.

Figure A.4: Motivating Evidence

(A) Distribution of Cuisine



(B) Distribution of HHI



Note: Panel (A) pie chart shows the distribution of cuisines among kitchens in our data. The cuisine variable is constructed based on kitchen features, dish features, and hometown information of the kitchen owners. Panel (B) shows the distribution of HHI constructed by seeing each consumer as a market. $HHI_i = \sum_j ms_{ij}^2 \cdot 10000$, where ms_{ij} is the market share of order from kitchen j among all of consumer i's orders.

B Permutation Test

In this section, we provide the detailed construction of the permutation test.[36] The null hypothesis for the permutation test is that consumers have no state-dependent preferences, and have an i.i.d fixed (unknown) probability of choosing each option from the choice set. The first step to test for variety-seeking preferences in consumer i is to observe his/her order sequence and compute the switching probability from a kitchen on the platform, $\hat{P}^i(y_t \neq j | s_t = j \neq 0)$, where s_t represents the previous period choice and y_t is the current period choice. The second step is to compute this switching probability for each permutation of the observed order sequence: each of these permutations is equally likely because consumer i's probability of choosing each option is fixed under the null hypothesis. Note that the implicit assumption here is a stochastic choice model: for each period, the probability of choosing each option is i.i.d across different periods. The set of switching probabilities and their respective relative frequencies computed in the second step constitutes the sampling distribution of switching probabilities under the null hypothesis, conditional on the observed number of orders for each option. We use this distribution for statistical testing.

Let $\mathbf{Y}$ be an order sequence from consumer i, and $\mathbf{N_i}$ denote the number of orders for all chosen kitchens by consumer i. The null hypothesis is that the choice probabilities are i.i.d across periods with $P(Y_t = j) = p_i^j$. This implies that, conditional on the number of orders for each kitchen, $\mathbf{N_i}(\mathbf{Y}) = \mathbf{n_i}$, each permutation is equally likely. The variety-seeking hypothesis predicts that the switching probability $\hat{P}^i(y_t \neq s_t | s_t \neq 0)$ will be significantly larger than what one would expect by chance. Let $\hat{P}(\mathbf{Y})$ denote this switching probability for sequence $\mathbf{Y}$. For

[36] Miller and Sanjurjo (2018) perform a similar permutation test as a nonparametric robustness test that is by construction invulnerable to conditional probability bias from small samples. Miller et al. (2014) showed that the runs and serial correlation tests, along with the conditional probability test, all amount roughly to the same test, and moreover, that they are not sufficiently powered. Traditional research on state dependence uses the runs tests to test for state dependence, such as the binomial runs test and multinomial runs test discussed in Bass et al. (1984). The permutation test is equivalent to the runs test in binomial choices. In multinomial choices, the multinomial runs test needs an extra assumption that the probability of choosing a brand on each choice occasion is equal to the observed empirical frequency of choosing the brand, whereas in the permutation test, we do not need to make the extra assumption and only need the choice probabilities to be i.i.d across periods.

an observed sequence $\mathbf{y}$, with $\mathbf{N_i}(\mathbf{y}) = \mathbf{n_i}$ orders for each chosen kitchen on the platform, to test the null hypothesis at the α level, we check if $\hat{P}(\mathbf{y}) \geq c_{\alpha,\mathbf{n_i}}$, where the critical value $c_{\alpha,\mathbf{n_i}}$ is defined as the smallest c such that $\mathbb{P}(P(\mathbf{Y}) \geq c|H_0, \mathbf{N_i}(\mathbf{y}) = \mathbf{n_i}) \leq \alpha$, and the distribution $\mathbb{P}(P(\mathbf{Y}) \geq c|H_0, \mathbf{N_i}(\mathbf{y}) = \mathbf{n_i})$ is generated from the sampling distribution from permutations. For the quantity $\mathbb{P}(P(\mathbf{Y}) \geq c_{\alpha,\mathbf{n_i}}|H_0, \mathbf{N_i}(\mathbf{Y}) = \mathbf{n_i})$, it may be the case that, for some c^*, it is strictly greater than α for $c \leq c^*$. In this case, for any sequence with $\mathbf{N_i}(\mathbf{y}) = \mathbf{n_i}$, one cannot reject H_0 at an α level of significance.[37] From the ex-ante perspective, a test of variety seeking at the α level of significance consists of a family of such critical values $\{c_{\alpha,\mathbf{n_i}}\}$. It follows immediately that $\mathbb{P}(reject|H_0) \leq \alpha$ because $\mathbb{P}(reject|H_0) = \sum_{\{\mathbf{n_i}\}} \mathbb{P}(P(\mathbf{Y}) \geq c_{\alpha,\mathbf{n_i}}|H_0, N_i^j(\mathbf{Y}) = \mathbf{n_i})\mathbb{P}(\mathbf{N_i}(\mathbf{Y}) = \mathbf{n_i}|H_0) \leq \alpha$. Last, for arbitrary test statistic $T(\mathbf{Y})$, the fact that the distribution of $\mathbf{Y}$ is exchangeable conditional on $\mathbf{N_i}(\mathbf{Y}) = \mathbf{n_i}$ means that $\mathbb{P}(P(\mathbf{Y}) \geq c_{\alpha,\mathbf{n_i}}|H_0, \mathbf{N_i}(\mathbf{Y}) = \mathbf{n_i})$ can be approximated to arbitrary precision with Monte Carlo permutations of the sequence $\mathbf{y}$ (Ernst, 2004; Good, 2013).

C More Robustness Checks

In this section, we provide more robustness checks on spurious state dependence and alternative model specifications.

C.1 Spurious State Dependence

C.1.1 Misspecified Heterogeneous Preferences

To confirm that we adequately control for heterogeneity, we estimate the model based on a reshuffled sequence of brand choices for each individual consumer. We no longer find evidence of state dependence with the reshuffled sequence, suggesting the state dependence we observe is not pure spurious state dependence resulting from misspecified heterogeneity. [38] We perform

[37] Examples of such choice sequences are AAAA and ABCD, where $\mathbb{P}(P(\mathbf{Y}) \geq c_{\alpha,\mathbf{n_i}}|H_0, \mathbf{N_i}(\mathbf{Y}) = \mathbf{n_i})$ always equal to zero and one, respectively. In such cases, we are not able to reject the null hypothesis that there is no state dependence and switching is simply due to randomness in choices.

[38] Dubé et al. (2010) adopted a similar method of robustness check for spurious state dependence.

an additional test for the possibility that the state dependence resulted from a misspecification of the distribution of heterogeneity. Suppose there is no structural state dependence and that the assumed functional form of heterogeneity causes the spurious state dependence. Then, if we randomly reshuffle the order of choice situations, the coefficient on the lagged choice will not change and still provide misleading evidence of spurious state dependence. The 2.5 percentile, median, and 97.5 percentile values of the log marginal likelihoods for the main model, which we fitted to 30 randomly reshuffled purchase sequences, are -304072, -303960, and -303762. A 95% interval of the log marginal likelihood based on the reshuffled purchase sequences contains the log marginal likelihood pertaining to the model that does not include a state dependence term (-303835). Furthermore, the 95% interval is strictly below the marginal likelihood pertaining to the model that includes a state dependence term based on the real choice sequence (-297610). We thus find additional evidence against spurious state dependence from a misspecified preference heterogeneity.

C.1.2 Negative Autocorrelation

One potential source of spurious state dependence is from negative autocorrelation in the unobserved taste shocks. If the choice model errors are negatively autocorrelated, a past purchase will proxy for a large past and hence a small current random utility draw.[39] The policy implications for the two models will be different since in the autocorrelated error case sellers can not use marketing instruments to influence future demand.[40] We test this by interacting the lagged choice with a dummy variable of whether the consumer's favorite kitchen was available in the previous period, where the favorite kitchen is defined by the one that is ordered most during the time window we study. Specifically, the indirect utility that consumer i derives from kitchen j

[39]Unlike the positive autocorrelation case, which could capture the seasonality or temporal cravings in preference for food, this is more of a mathematical possibility than an intuitive preference possibility.

[40]From the firm's point of view, autocorrelation is equivalent to a more dispersed distribution of unobserved heterogeneity, whereas in the variety-seeking case, prices affect both current period demand and future distribution of consumer states, which gives rise to a nontrivial dynamic pricing problem. If sellers lower their current price, their own consumer share will increase, and sellers will face lower and more elastic demand in the next period.

on platform in period t is

$$u_{ijt} = \beta_{ij} - \alpha_i p_{jt} + \eta r_{ijt} + \gamma_i \mathbb{1}(s_{it} = j) + \gamma^F \mathbb{1}(s_{it} = j)\mathbb{1}(k_i \notin \mathcal{J}_{i(t-1)}) + \psi \cdot X_{ijt} + \epsilon_{ijt}, \quad (21)$$

where k_i denotes consumer i's favorite kitchen.[41] Under variety seeking, whether the previous period choice was ordered when the consumer's favorite kitchen was available should not affect the state-dependent preferences. However, when ϵ_{ijt} are negatively autocorrelated, kitchens ordered when the favorite kitchen was not available will have a lower $\epsilon_{ij(t-1)}$, which implies a higher ϵ_{ijt}. This should decrease the switching probability in the current period, i.e., consumers should switch less away from kitchens that are ordered when the favorite kitchen was not available.($\gamma^F > 0$) Column (3) of Table C.1 presents the result of this alternative specification. We find that γ^F is positive but insignificant, which implies consumers do not switch less away from kitchens ordered when the favorite kitchen was not available. This could result from the fact that our sample is mainly composed of the frequent users who ordered at least once a week, which limits the variation in the interpurchase time span of consumption.

C.2 Initial Conditions Bias

For each consumer choice panel, we drop the first observation of purchase and keep the choice as the state variable for the next purchase occasion. The implicit assumption behind this practice is that the initial state is exogenous and independent of preference. Simonov et al. (2020) discuss the potential initial conditions bias in demand estimation with state dependence and shows that this assumption could overestimate loyalty in small samples. On the one hand, our panel is relatively long with 40 choices per consumer on average, which reduces the small-sample bias. On the other hand, in the case that consumers are more likely to order the product that they prefer more in the initial period, treating the initial state as exogenous will attribute preference to loyalty and generate a conservative estimate of the variety-seeking preferences.

For orders put through channels other than the homepage listing, such as the discover

[41] The favorite kitchen is defined as the kitchen with the highest choice probability conditional on availability.

channel or the favorite kitchen folder. We drop the specific period's choice situation from the consumer panel and use the observed choice to update the state variable for the next period. Following similar arguments as in the initial conditions bias discussion above, if we assume that in these orders placed through channels other than the rank listing, consumers are more likely to order the product that they prefer more, then treating the state from these periods as exogenous will generate a conservative estimate of variety-seeking preferences. The same is true for orders placed on pages other than the first page, for which we drop the choice situations but update the state variable for the next period according to the actual order. When consumers search and order beyond the first page, it is very likely that consumers are looking for a specific niche kitchen that they like. The same analysis of the direction of bias by treating states from these periods as exogenous applies. Moreover, orders of kitchens on the first page through the observed ranking list take 83% of the total number of orders, so the scope of this potential problem is limited.

C.3 Dish-Level Variety Seeking

In the baseline model we do not consider the dish-level variety-seeking preferences, while in reality, consumers could be switching across different dishes within the same kitchen. To test for this possibility, we add an interaction term of the last kitchen choice with the menu size of the kitchen. Specifically, the indirect utility that consumer i derives from kitchen j on platform in period t is

$$u_{ijt} = \beta_{ij} - \alpha_i p_{jt} + \eta r_{ijt} + \gamma_i \mathbb{1}(s_{it} = j) + \gamma^M \mathbb{1}(s_{it} = j) M_{jt} + \psi \cdot X_{ijt} + \epsilon_{ijt}, \tag{22}$$

where M_{jt} denotes the menu size of kitchen j in period t. The utility from choosing the outside option is $u_{i0t} = \epsilon_{i0t}$. If consumers are switching across different dishes within the same kitchen, we should expect consumers to switch less away from kitchens with a larger menu size, i.e., $\gamma^M < 0$. Column (1) of Table C.1 presents the result of this alternative specification. We find that γ^M is not significantly different from zero, which implies consumers do not switch less

Table C.1: Robustness Check

		(1) Menu Size	(2) Interpurchase Time Span	(3) Negative Autocorrelation
State Dependence	μ_γ	-0.197***	-0.204***	-0.243***
–Kitchen level (γ_i)		(0.033)	(0.051)	(0.067)
	σ_γ	0.642***	0.501***	0.521***
		(0.050)	(0.067)	(0.058)
Price (α_i)	μ_α	-0.028***	-0.027***	-0.017***
		(0.007)	(0.001)	(0.001)
	σ_α	0.012***	0.018***	0.011***
		(0.001)	(0.004)	(0.001)
Rank	η	-0.048***	-0.037***	-0.051***
		(0.006)	(0.005)	(0.005)
Menu Size	γ^M	-0.003		
		(0.007)		
Interpurchase Time Span	γ^T		0.018*	
			(0.011)	
Favorite Seller Unavailable in Previous Period	γ^F			0.008
				(0.006)

Note: The table shows estimation results after controlling for cuisine, day of the week, week of the month, holiday, menu size, distance, rating, monthly sales, and kitchen fixed effects. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Column (1) tests for dish-level variety seeking. Column (2) tests for the effect of interpurchase time span on the variety-seeking level. Column (3) tests for robustness to negative autocorrelation.

away from kitchens with larger menu sizes.[42] This could happen when consumers have one or two favorite dishes that they always order from a specific kitchen.

C.4 Interpurchase Time Span

When constructing the choice panel, we drop the days when consumers did not use the platform, which makes the interpurchase time spans different across different periods in terms of calendar days. Intuitively, calendar day should be a natural period of satiation over tastes, and the variety-seeking preferences could decrease with the interpurchase time span across platform

[42]We do not model variety seeking on the dish level mainly for tractability. Future research could study state dependence on a finer level, including dishes, ingredients, recipes, and cooking methods. Natural language processing and graph classification could be helpful methods for such analysis.

usage. To test this interpurchase time span effect, we add an interaction term of the last kitchen choice with the number of calendar days from the last consumption. Specifically, the indirect utility that consumer i derives from kitchen j on platform in period t is

$$u_{ijt} = \beta_{ij} - \alpha_i p_{jt} + \eta r_{ijt} + \gamma_i \mathbb{1}(s_{it} = j) + \gamma^T \mathbb{1}(s_{it} = j) T_{ijt} + \psi \cdot X_{ijt} + \epsilon_{ijt}, \tag{23}$$

where T_{ijt} denotes the interpurchase time span of consumer i with kitchen j until period t. The utility from choosing the outside option is $u_{i0t} = \epsilon_{i0t}$. If the state-dependent preferences decrease with interpurchase time span, we should have $\gamma^T > 0$. Column (2) of Table C.1 presents the result of this alternative specification. We find that γ^T is positive but insignificant at the 5% level, which implies consumers do not switch less away from kitchens with longer interpurchase span. This can result from the fact that our sample is mainly composed of the frequent users who ordered at least once a week, which limits the variation in the interpurchase time span of consumption.

C.5 Model Fit

We examine model fit by comparing several key predictions to their empirical counterparts. We split the sample of 6,629 consumers randomly into a 80% training sample (5,303 consumers) and a 20% holdout sample (1,326 consumers). We estimate the baseline model on the training sample and use the estimated model to predict several key moments in the holdout sample. Table C.2 presents the simulated predictions and the empirical counterparts from the holdout sample. We evaluate model's out-of-sample fit in terms of purchase probability, switching probability, purchase position, and purchase probabilities by position. We find that the model predicts the holdout sample data reasonably well in terms of the key moments.

Table C.2: Model Out-of-Sample Fit

I. Key Moments		
	Data	Predicted
Purchase Probability	0.473	0.481
Purchase Position	3.464	3.457
Switching Probability	0.814	0.821
II. Purchase Probabilities by Position		
Position	Data	Predicted
1	0.145	0.131
2	0.118	0.127
3	0.047	0.050
4	0.034	0.035
5	0.031	0.033
6	0.027	0.026
7	0.025	0.024
8	0.022	0.020
9	0.019	0.017
10	0.017	0.016

Note: The table reports results from an out-of-sample validation exercise. We split the sample of 6,629 consumers randomly into an 80% training sample (5,303 consumers) and a 20% holdout sample (1,326 consumers). We estimate the baseline model on the training sample and use the estimated model to predict several key moments in the holdout sample.

D Counterfactual Analysis Details

D.1 Optimality of Utility-based Ranking

We prove the optimality of ranking kitchens in decreasing order of expected utility in this section. The ex-ante consumer welfare is

$$E[CS] = -\frac{1}{\alpha}\left[ln\sum_{j\in\mathcal{J}} exp(v_j - c_j)\right]$$

where α is the price coefficient, v_j represents the deterministic part of utility excluding the position cost, and $c_j = \eta \cdot r_j, r_j = 1, 2, ..., J$ represents the position cost. Then the optimal ranking should maximize $E[CS]$ by assigning $c_1, c_2, ..., c_J$ to kitchens in the choice set, represented by $v_1, v_2, ..., v_J$. Without loss of generality, we assume that $v_1 > v_2 > ... > v_J, c_1 < c_2 < ... < c_J$.

Lemma: Ranking kitchens in the decreasing order of expected utility v_j optimizes ex-ante consumer welfare, given $v_1, v_2, ..., v_J, c_1, c_2, ..., c_J$.

Proof: Consider the two product case where $v_1 > v_2, c_1 < c_2$, we want to prove that the expected welfare from $\{v_1 - c_1, v_2 - c_2\}$ is greater than $\{v_1 - c_2, v_2 - c_1\}$. Denote $x_1 = v_1 - c_1, x_2 = v_2 - c_2, \tilde{x}_1 = min\{v_1 - c_2, v_2 - c_1\}, \tilde{x}_2 = max\{v_1 - c_2, v_2 - c_1\}$, then $x_1 + x_2 = \tilde{x}_1 + \tilde{x}_2, x_1 < \tilde{x}_1 < \tilde{x}_2 < x_2$. Because the exponential function is increasing and convex, we have

$$\begin{aligned} exp(x_2) - exp(\tilde{x}_2) &> exp(\tilde{x}_2) \cdot (x_2 - \tilde{x}_2) \\ &> exp(\tilde{x}_1) \cdot (\tilde{x_1} - x_1) \\ &> exp(\tilde{x}_1) - exp(x_1) \end{aligned}$$

Thus $E[CS_{\{v_1-c_1, v_2-c_2\}}] > E[CS_{\{v_1-c_2, v_2-c_1\}}]$.

Then, given a general number of products J, we want to prove that the optimal ranking is $\{v_1 - c_1, v_2 - c_2, ..., v_J - c_J\}$. Starting from any ranking position assignment $R_1 = \{..., v_j - c_j, ..., v_k - c_k, ...\}$ where $\exists j, k, s.t., v_k > v_j, c_k > c_j$, we show that this assignment is dominated by reversing the position of the two kitchens while keeping the other positions fixed $R_2\{..., v_j - c_k, ..., v_k - c_j, ...\}$.

$$\begin{aligned} E[CS_{R_1}] = &\ -\tfrac{1}{\alpha}\left[ln(exp(v_j - c_j) + exp(v_k - c_k) + C)\right] \\ E[CS_{R_2}] = &\ -\tfrac{1}{\alpha}\left[ln(exp(v_j - c_k) + exp(v_k - c_j) + C)\right] \end{aligned}$$

where C denotes the sum of exponential functions for the other positions. Then from the proof in the case when $J = 2$, we have $E[CS_{R_1}] < E[CS_{R_2}]$. We can continue this process until there is no higher expected utility assigned to lower ranking positions, and thus we prove the optimality of expected utility based ranking.

Note that the utility-based ranking need not optimize the platform revenue, since the platform could promote lower utility but higher price kitchens to increase revenue. There is no simple index like in the utility case, and the ranking will face a tradeoff between price and

choice probability. Appendix 6 in Compiani et al. (2021) provides more details on the intuition of relevant tradeoffs of revenue-maximizing ranking algorithms. The crucial value in the optimality of ranking algorithms in this case is the difference between option j's revenue and the choice-probability weighted average revenue from all other options.[43]

D.2 Numerical Solution for Price Competition

We use numerical methods to solve for the equilibrium of the pricing game. We first discretize each axis of the state space using a finite number of points. We then form a grid representing the whole state space from the Cartesian product of these points, excluding the support where the sum of market shares exceeds one. For each point in the grid, we store the value and policy functions of each competitor in the computer memory. For states outside the grid, we calculate the value and policy function using polynomial interpolation.

To solve for the equilibrium, we employ the following algorithm, which is an adaptation of policy iteration applied to the case of the games: start with some initial guess of the strategy profile, $\sigma^0 = (\sigma_1^0, \sigma_2^0)$, and then proceed along with the following steps:

1. For the strategy profile σ^n, calculate the corresponding value functions for each of the J kitchens. These value functions are defined by the Bellman equation, where the right-hand side of the Bellman equation is maximized, given the current strategy profile of other kitchens. Update each kitchen's strategy using the Bellman equation. Denote the resulting new policies and value functions by σ_j^{n+1} and V_j^{n+1}.

2. If $n > 0$, check whether the value functions and policy functions satisfy the convergence criteria, $||V_j^n - V_j^{n-1}|| < \epsilon_V$ and $||\sigma_j^n - \sigma_j^{n-1}|| < \epsilon_\sigma$ for all kitchens j. If so, stop. If not, return to step 1.

[43]The industry ranking algorithm usually uses the multi-objective optimization model. Food Discovery with Uber Eats: Recommending for the Marketplace `https://eng.uber.com/uber-eats-recommending-marketplace/` In this project, we focus on documenting, measuring variety seeking, and showing its importance through counterfactual analysis, instead of providing the exact engineering algorithms.

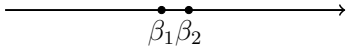
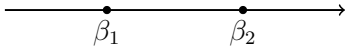
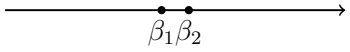
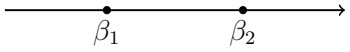
D.3 Intuition for Variety Seeking Effect on Price Competition

In this section, we provide more intuition on the effect of consumers' variety-seeking preferences on price competition. Seetharaman and Che (2009) find that variety seeking softens price competition and Sajeesh and Raju (2010) find that variety seeking intensifies price competition. Zeithammer and Thomadsen (2013) shows that when products are similar and variety seeking is strong, variety seeking softens price competition, whereas when products are dissimilar and variety seeking is weak, variety seeking intensifies price competition. This seems in conflict with the traditional wisdom from industrial organization that differentiation softens price competition. Here, we discuss how we can understand this result intuitively and how it connects with our empirical findings. Table D.1 shows the product differentiation level without the variety-seeking effect. The columns correspond to two cases when products are similar, and variety seeking is relatively strong, and when products are dissimilar and variety seeking is relatively weak. The rows are products' values from consumers who purchased product 1, and 2 in the previous period, respectively.

Table D.2 shows the product differentiation level with the variety seeking effect. For consumers who purchased the inferior product in the last period, variety seeking decreases the product value further and increases the differentiation level. For consumers who purchased the superior product in the last period, variety seeking decreases the product value of the superior product, and the effect on the differentiation level depends on the relative magnitude of static differentiation from product value, and the dynamic differentiation from state dependence. When products are similar and variety seeking is strong, variety seeking increases overall differentiation and softens price competition, whereas when products are dissimilar and variety seeking is weak, variety seeking decreases overall differentiation and intensifies price competition.[44] Our demand estimates suggest that we are in a case where the top two popular groups

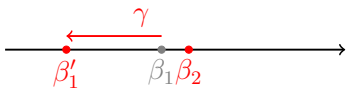
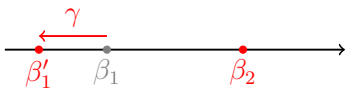
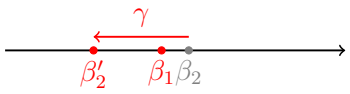
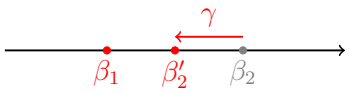
[44]Variety seeking intensifies price competition when products are dissimilar, and there is one big popular seller maintaining its own consumers who are choosing between consuming the superior product consecutively and switching towards the inferior product. This decomposition is also helpful for understanding the nonmonotonic relationship between state dependence level and price competition level. Holding fixed the static product values, there is a range where state dependence first decreases and then increases the product differentiation level.

Table D.1: Product Differentiation without Variety-Seeking Effect

	Similar, strong VS	Dissimilar, weak VS
Seller 1's consumer:		
Seller 2's consumer:		

Note: The table shows product differentiation level without the variety seeking effect. The columns correspond to two cases when products are similar and variety seeking is relatively strong, and when products are dissimilar and variety seeking is relatively weak. The rows are products' values from consumers who purchased product 1, and 2 in the previous period, respectively.

Table D.2: Product Differentiation with Variety Seeking Effect

	Similar, strong VS	Dissimilar, weak VS
Seller 1's consumer:	 Differentiation ↑	 Differentiation ↑
Seller 2's consumer:	 Differentiation ↑	 Differentiation ↓

Note: The table shows product differentiation level with the variety seeking effect. The columns correspond to two cases when products are similar and variety seeking is relatively strong, and when products are dissimilar and variety seeking is relatively weak. The rows are products' values from consumers who purchased product 1, and 2 in the previous period, respectively. Arrows show the differentiation level change in each case.

are relatively similar, and consumers' variety-seeking preferences are relatively strong, so variety seeking softens price competition within a range around our demand estimate level.

E Industry Practice

In this section, we discuss the current targeted marketing practice in the industry, and the potential improvement by taking into account consumer variety-seeking preferences. We find

that consumers' variety-seeking preferences are not accounted for in current industry practice of targeted pricing and are likely not emphasized enough in ranking and recommendation systems on several common food delivery platforms.

E.1 Targeted Pricing

Current industry practice in pricing is more observable than ranking and recommendation systems on food delivery platforms. We find several common platforms' promotion activities facing merchants on their websites. Uber Eats allows merchants to target promotions to either all customers in their delivery zone or new customers who have not ordered from their restaurant. DoorDash's 'Order Again and Save' promotion is for customer retention and allows merchants to pay $0.99 fee per order to target their existing DoorDash consumers who haven't ordered from them in over 45 days. Similarly, Grubhub's promotion for past customers is also focused on past customer retention, and it only allows merchants to target consumers who have not ordered in the last 90 days to order again. None of the popular food delivery platforms allows targeted pricing to old customers with recent orders. Table E.1 provides a summary of the current merchant-level promotions allowed on the popular food delivery platforms.[45]

E.2 Personalized Recommendation

Besides ranking and pricing, our findings also have important implications for the design of personalized recommendations. Figure E.1 shows the current common practice of personalized recommendations on several popular food delivery platforms. The recommendation sections, such as UberEats's 'Order Again', Grubhub's 'Reorder your favorites', and DoorDash's 'Your Favorites, all understand and emphasize the importance of consumers' persistent heterogeneous

[45]Information is from the platforms' public websites for merchant promotions support. Uber Eats: https://help.uber.com/merchants-and-restaurants/article/uber-eats-manager-marketing-tool-faqs?nodeId=a9838c22-2065-43e1-85c2-cbca0d3e9802 DoorDash: https://help.doordash.com/merchants/s/article/Getting-started-with-DoorDash-Promotions?language=en_US#What%20kinds%20of%20promotions%20does%20DoorDash%20offer? https://help.doordash.com/merchants/s/article/Order-Again-and-Save-Promo-FAQ?language=en_US Grubhub: https://get.grubhub.com/help-center/how-to-run-promotions/ https://lp.grubhub.com/restaurant/smart-promotions/

Table E.1: Current Targeted Pricing Available for Merchants on Popular Food Delivery Platforms

Platform	Current Consumer Types for Merchant Targeted Promotions
Uber Eats	New consumers
	All consumers in the restaurant's delivery zone
DoorDash	New consumers
	Consumers who haven't ordered from the restaurant in over 45 days
	All consumers
Grubhub	New consumers
	Consumers who haven't ordered in the last 90 days
	All consumers

Note: The table summarizes the current promotions allowed on popular food delivery platforms. Information comes from the platforms' websites for merchants. None of the popular platforms allows for promotions targeting existing consumers who ordered recently.

preferences, but it does not seem like they consider consumer variety-seeking preferences by avoiding recent orders.

Koren et al. (2009) discuss several techniques for recommendation systems in the context of Netflix movie recommendations. Such systems are particularly useful for entertainment products such as movies, music, and TV shows. Among the models discussed, the only model that captures dynamics in demand is the temporal dynamics model, which accounts for the time-drifting nature of user preference by allowing preference parameters as a function of time. But the specific model captures the changes in persistent preference over time,[46] describes this with an example:"..., a fan of the psychological thriller genre might become a fan of crime dramas a year later. Similarly, humans change their perception of certain actors and directors", which is different from the high-frequency dynamics driven by variety-seeking preferences.

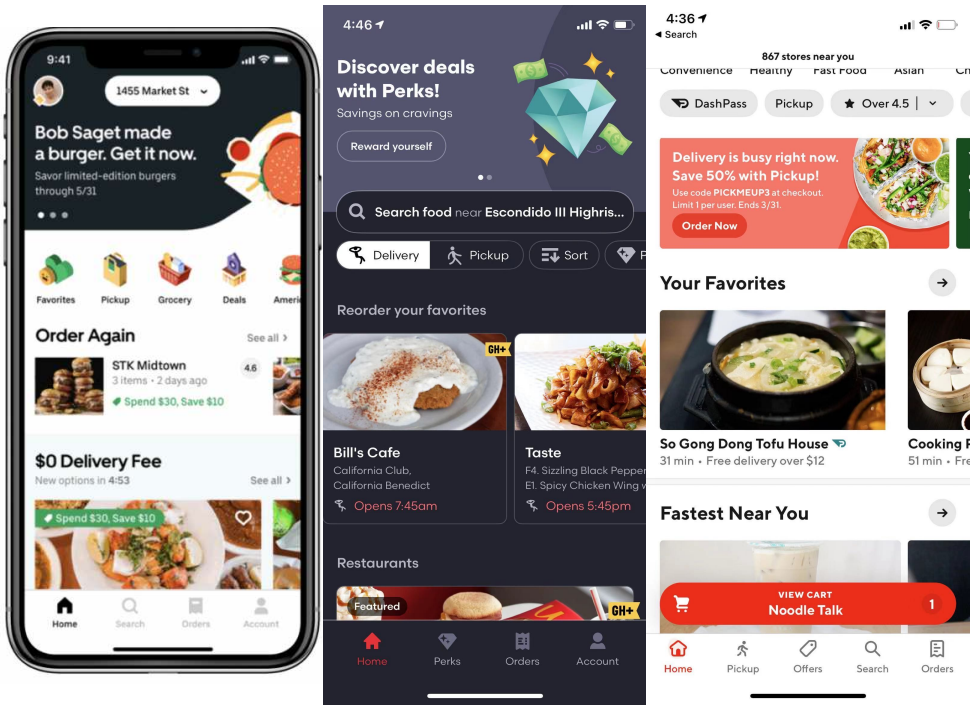
Uber Eats posts research results on Uber Engineering blogs, and based on the recent reports, variety-seeking dynamics are not captured in their recommendation or ranking algorithm. Jain et al. (2019) [47] discuss that Uber Eats uses the graph learning method GraphSAGE developed by Hamilton et al. (2017) to power their recommendation system. In this method, the aggregators are symmetric(permutation invariant). [48] and a random permutation of the users' orders does

[46]Source: Koren et al. (2009)

[47]Food Discovery with Uber Eats: Using Graph Learning to Power Recommendations `https://eng.uber.com/uber-eats-graph-learning/`

[48]The method is based on a graph where nodes are users and sellers, and edges are weighted by the historical

Figure E.1: Recommendations on Popular Food Delivery Platforms



Note: The graph shows screenshots of the homepage recommendation from popular platforms, including Uber Eats, Grubhub, and DoorDash. The frequently ordered kitchens are recommended as 'Order Again' and 'Your Favorites' without punishment for the recency of orders from the recommended kitchens.

not affect the graph and thus recommendation, thus the high frequency dynamics induced by variety-seeking preferences are not included in this method or widely used matrix factorization and collaborative filtering methods for recommendation systems. Meanwhile, methods that process inputs in a sequential manner, such as LSTM, have limitations in scalability.

E.3 Upper Confidence Bound Algorithm

Currently one of the most commonly used optimal solutions for multi-armed bandit problems is the upper confidence bound (UCB) algorithm.[49] One general form of the method is to choose

number of orders. For example, consumer 1 who ordered CBAAA and consumer 2 who ordered ABACA will generate the same graph and the method will predict that they will have similar order patterns.

[49]This method is currently used by Uber Eats. "..., we applied the upper confidence bound (UCB, one of the methodologies of MAB) approach to facilitate the exploration of new restaurants or restaurants with low impressions. We calculate a UCB score for each restaurants based on metrics such as its historic impression number, total click number, and boosting factor. The UCB score, along with other objectives discussed above, decides the ranking order among restaurants. A new restaurant will have a relatively high UCB score initially and hence rank highly, increasing its exposure. As the new restaurant gathers other impressions, the UCB score will smoothly decrease and gradually transfer ranking weight back to other objective scores such as relevance."

the option with the optimal UCB index, UCB_i, which is defined as follows:

$$UCB_i = \hat{\mu}_i + f(T_i) \tag{24}$$

where $\hat{\mu}_i$ is the empirical mean of former experience. $f(T_i)$ is a decreasing function of the sample times the learner has observed for arm i. This index illustrates the explore vs exploit tradeoff, where the first term $\hat{\mu}_i$ captures the exploit part and the second term $f(T_i)$ represents potential reward from exploration. The explore vs exploit tradeoff implied by this solution suggests less switching among familiar arms and more switching towards new arms, which corresponds to the two tests we perform in Section 4.5.1 and Section 4.5.2. Note that the explore expression is different from variety seeking in the sense that it is based on cumulative experience rather than the recency of experience.

Wang et al. (2018): Food Discovery with Uber Eats: Recommending for the Marketplace `https://eng.uber.com/uber-eats-recommending-marketplace/`