Choice Frictions in Large Assortments

Olivia R. Natan*

February 15, 2023

Abstract

This paper studies how the growth and evolution of product assortments impact consumer adoption, churn, and long run consumption. Most economic theories of product variety and the value of platforms suggest consumers at least weakly prefer larger product assortments. In contrast, the psychological literature on the phenomenon of choice overload finds that larger assortments overwhelm consumers with decision costs or induce more regret. I provide empirical evidence of how the size and contents of product assortments impact consumers over their lifetime on an online food delivery platform. I find that assortment expansion increases the acquisition of new consumers but reduces the frequency of consumption among consumers who remain on the platform. I rationalize these effects via a model of costly attention and choice under limited information. Counterfactual exercises show that targeting choice set reductions can improve revenue among existing customers.

*University of California, Berkeley Haas School of Business; olivia.natan@haas.berkeley.edu. I am grateful to Jean-Pierre Dubé, Pradeep Chintagunta, Anita Rao, and Ali Hortaçsu for their advice and support. Winnie Xiao provided excellent research assistance. Thanks go to Junhong Chu, Ali Goli, Joonhwi Joo, Uyen Tran, Walter Zhang, Stephan Seiler, Raluca Ursu, Julie Holland Mortimer, Connor Dowd, and participants at various seminars for their thoughtful comments. I thank the data provider; all errors are my own.

1 Introduction

Consumers have access to more product variety now than ever before. Both online and offline, retailers offer consumers greater product variety through larger and more varied assortments.[1] Whether this increased variety is inherently beneficial to consumers is unclear. Psychologists have documented that presenting more choice alternatives to consumers may lower the likelihood of purchase (Iyengar and Lepper, 2000). Choice overload—the negative effect of additional choice alternatives on consumer purchase and satisfaction—has been sometimes demonstrated in small assortments both in the lab and in small field tests (Chernev et al., 2012; Scheibehenne et al., 2010). Despite this evidence, economists approach product variety by using models that assume 'more is better' for individual consumers. Widely used models of consumer demand preclude the possibility that consumers may prefer fewer choice alternatives.

An empirical literature on retail assortments has found mixed results on the impact of assortment size on category sales (Drèze et al., 1994; Broniarczyk et al., 1998), interpurchase time (Borle et al., 2005), and store choice (Briesch et al., 2009). These studies primarily focused on changes to the store-level variety and a single purchase occasion. Changes to product assortments may affect customers differently, depending on category familiarity, their taste for variety, and the match between new products and their tastes. Optimal assortment strategy will reflect this consumer heterogeneity.

In this paper, I examine empirically how larger product assortments affect individual consumption. I document the dynamic impact of changing product assortments on consumer acquisition, short-term retention, and long-term consumption frequency using quasi-experimental variation. I find that more variety is detrimental to the purchase frequency of existing customers, though consumer adoption into the category grows due to greater product variety. Consumers engage in longer search when they face larger assortments, and experiencing these higher search costs reduces their future purchase frequency. To separate the cost of combing through variety from possible changes to consumer match value, I construct and estimate a model of attention allocation where consumers' beliefs and attention costs depend on assortment size. I show that consumers' expectations about product valuations change, but information costs are not directly affected. I demonstrate how only individually-targeted assortment reductions can improve sales among existing consumers.

The setting for this paper is an online restaurant-to-consumer delivery platform where the

[1] The average grocery retailer carries nearly 50,000 SKUs in 2008, up from 9,000 in 1975. (Food Marketing Institute). Spotify offers over 40 million songs (Aguiar et al., 2021).

assortment varies across both consumers and time. In this context, product variety is measured as the number of restaurants and breadth of cuisines that deliver via the platform to a consumer's location. As the platform expands over time, consumers face growing numbers of restaurants that will deliver to their location. Unlike many offline retail contexts, the variation in assortment size is observed instead of inferred from purchase data. The impact of assortment growth in online restaurant-to-consumer delivery markets can be isolated from other confounding factors, since prices are typically fixed and new consumers have experience in similar categories.

Consumers adopt the platform at higher rates in neighborhoods with larger local assortments of restaurants. Conditional on adoption, I observe higher consumer spending in neighborhoods with larger local assortments. Since households likely choose neighborhoods based in part on local amenities including restaurants, I cannot use this raw correlation to determine the causal effect of variety. To account for the endogeneity generated by household location choices, I use changes to consumers' choice sets across time and geographic space to identify the causal effect of the size of the assortment. This approach compares the within-household variation in choice sets and purchase behavior among households in the same neighborhood who receive different assortments. I find that more restaurants increase the number of new consumers who come to the platform but reduce the order frequency for existing consumers.

Next, in order to quantify the mechanism by which assortment expansion negatively impacts existing customers, I construct and estimate a structural model of consumer attention and demand where beliefs about match value and information costs may vary with the assortment size. This model—rational inattention—proposes that consumers select how much costly information to acquire about choice alternatives before making a discrete purchase decision. The model, which builds off of Joo (2022), nests a test of how larger assortments increase information costs separately from consumers' beliefs, unlike ad hoc adjustments to standard discrete choice models. I reject that the cost of a unit of information changes as the assortment grows. Instead, I find that inter-purchase time is increasing in the assortment size because consumers' expectations of untried products are lower as the assortment grows.

Finally, I use the structural model to test how much the platform can offset the downsides of assortment expansion by offering personalized choice sets. Revenues can be improved by offering different types of targeted assortment reductions to different consumers. Testing counterfactual assortments is necessary to understand two countervailing forces in demand for heterogeneous products: (1) the negative effect of larger assortments on consumer expectations and (2) the improved possible match value generated by a larger assortment. I consider assortment restrictions that hold the supply and contents of restaurant fixed. For a potential restricted size of the assortment, I target the contents based on several metrics (e.g., probability of purchase, expected platform revenue). After simulating consumer choices and platform revenue under proposed assortment reductions, I compare revenue across both the magnitude of the assortment reduction and across targeting methods. I find that the platform can improve weekly purchase frequency up to 22% among existing customers by offering assortment reductions that target based on consumer preferences.

This paper relates closely to the marketing literature on product assortments, including work considering the possibility of making assortments strategically smaller. Much attention has been devoted to assortments at the store, category, and product line levels in grocery retailing (Briesch et al., 2009; Draganska and Jain, 2005; Drèze et al., 1994; Broniarczyk et al., 1998; Borle et al., 2005; Boatwright and Nunes, 2001).[2] I add to this stream of literature by measuring the effect of assortment size on individual consumers in platform adoption and repeat-purchase settings. Existing empirical work finds mixed results—in some contexts, removing low-selling items improves in-store sales, while in others, more products on the shelf yield higher sales.

Consumer behavior research has documented robust instances of choice overload, starting with Iyengar and Lepper (2000). These papers (reviewed in Chernev et al. (2012)) find that showing consumers a larger variety of products induces more interest in browsing the products, but fewer overall purchases. However, Scheibehenne et al. (2010) find that the average empirical effect is close to zero and depends on the context. I build on these small-scale, primarily lab-based, empirical findings by providing evidence for choice overload effects in an empirical setting with large choice sets, real consumption choices, and repeated consumption. Because of the richness of my empirical setting, I test how the cost of larger assortments differs across consumers and choice contexts.

Choice overload can be rationalized by several economic theories.[3] Consumers may draw inferences about product quality (Kamenica, 2008) or their utility net of search costs (Kuksov and Villas-Boas, 2010) from the length of the product line. The latter mechanism is exacerbated when firms are incentivized to garble the information presented to their customers (Nocke and Rey, 2022).

This paper also relates to the extensive economic literature on product variety. Theoretical work

[2]The structure of the assortment, such as shelf space allocation offline and product organization, can also influence consumer perceptions about the assortment's variety (Eisend, 2014; Kahn and Wansink, 2004). While changes to assortment size will be considered here, the fine-tuned adjustment of shelf facings is beyond the scope of this paper.

[3]Variety may negatively impact consumption levels for other reasons not explored by this paper. Examples include matching markets (Halaburda et al., 2018), product level direct network effects (Casadesus-Masanell and Halaburda, 2014), or consumer learning (Kim, 2021).

on product variety models consumers who receive utility directly from variety when they purchase a basket of goods (Dixit and Stiglitz, 1977; Bronnenberg, 2015) or due to many goods meeting heterogeneous tastes better (Hotelling, 1929; Lancaster, 1975, 1990). Empirical approaches to product variety have used the latter—the characteristics approach—to measure whether markets provide socially optimal product variety, and they typically find positive returns to variety on market size (e.g., Berry and Waldfogel (1999); Berry et al. (2016); Quan and Williams (2018); Illanes and Moshary (2020)). However, these discrete choice models of demand suffer from a mechanical issue with the introduction of new products. Each new product generates a new characteristic from each products demand shock, which mechanically increases consumer welfare (Ackerberg and Rysman, 2005). This paper takes the rational inattention approach to demand (Matêjka and McKay, 2015; Caplin and Dean, 2015) and extends Joo (2022) to estimate the returns to product variety in a manner that breaks this mechanical connection.

Lastly, my work also relates to empirical work on platforms and retailers focused on how the size of the seller base impacts competitive and demand dynamics on platforms. Two closely related works, Li and Netessine (2020) and Farronato et al. (2022), find evidence for limited-to-no cross network effects in online platforms, while others document positive cross-network effects (Chu and Manchanda, 2016; Lin, 2017). Reshef (2022) uses similar data and identifying variation to study the impact of assortment changes on how platform sellers price differently under increased competition. The findings of that paper—that new entries benefit 'strong' incumbents and hurt 'weak' ones—is consistent with the findings in this paper. Ershov (2022) similarly looks at how changes to search frictions changes entry quality; he finds a reduction in search costs on a platform spurs entry of low-quality products. I contribute to this literature by studying (i) how entry of sellers differentially impacts individual consumers based what part of the consumer lifecycle they are in and (ii) how platforms can leverage the online nature of their business to offer an individually-targeted solution. In practice, retailers reduce the scope of their assortments—a practice known as SKU rationalization—to address operational concerns. Limiting the number of products stocked on shelves simplifies operations and stocking costs. However, results of this strategy has proven mixed in offline contexts (Borle et al., 2005; Boatwright and Nunes, 2001; Sloot et al., 2006). WalMart, for example, ultimately reversed course after trying a large-scale SKU rationalization by bringing back 8,500 SKUs to their stores.[4] This paper provides evidence suggesting another mechanism by which retailers can benefit from SKU rationalization. In particular, if reductions can be customized

[4]Source: https://retailwire.com/discussion/walmart-reverses-course-on-sku-rationalization/

to individual customers, retailers can improve retention via SKU rationalization.

The rest of this paper is structured as follows. Section 2 describes the relevant details on the market context and the data. Section 3 presents the research design, reduced form results, and robustness checks. Section 4 presents the demand model, estimates, and counterfactual results. In Section 5, I conclude and discuss possible extensions.

2 Context and Data

2.1 Context

This paper studies product variety in the U.S. restaurant-to-consumer food delivery market. Restaurant delivery is a large market, accounting for 8% of restaurant sales in 2018.[5] Restaurant food reaches consumers at home via two channels: direct purchase from restaurants (mostly over the phone) and online ordering/delivery platforms. The direct channel typically allows consumers to place orders by calling local restaurants and waiting for a restaurant-employed driver to delivery food to their location. Some large restaurant chains offer direct online ordering and payment.[6] Since 2010, many online platforms for food ordering and delivery have entered the market. In this channel, consumers may use a single website or mobile app to access many local restaurants. The platforms allow consumers to see and select menu options, place and pay for their order, and receive delivery through a single service.

The online restaurant ordering and delivery market is a practically and substantively useful lab for studying the effects of product variety and assortment on individual consumers. Past work on assortments has focused mostly on grocery retail, but retail assortment data often is missing detailed availability (for example, due to stock-outs). Online restaurant delivery platform assortments are observed exactly based on platform data. Moreover, unlike offline assortments or many online retail assortments, individuals face different choice sets at the same time based on their location, and frequent entry and exit leads to considerable intertemporal assortment variation. Like offline restaurant delivery, each restaurant may select which neighborhoods or addresses they will deliver to. Neighbors can face different platform choice sets as a result of delivery zones.

In addition to the observed and widespread variation in assortments in this market, prices are typically fixed,[7] and new consumers have experience in similar categories (e.g. ordering restaurant

[5] See: https://www.statista.com/statistics/1091419/food-service-sales-delivery-share-us/ Mar 1, 2022

[6] For example, see Domino's Pizza online ordering system. Domino's alone sold $9.8 Billion in delivery pizza in the US in 2018 (See: https://ir.dominos.com/static-files/593a1150-28b1-49a9-8263-88710bb1237a, March 1, 2022)

[7] I will discuss prices further in section 2.2. Changes to the price of a meal happen for several reasons. First, prices of a meal differ across restaurants, including within cuisine. An upscale sushi restaurant may charge more for a single meal than both a local pizzeria and a casual neighborhood sushi restaurant. Second, the meal price can differ within

food over the phone or in person). These context-specific factors allow me to rule out alternative explanations for why assortment expansion could be detrimental. Consumers who are new to this market are not new to the broader category of prepared food. Consumption through the platform may allow them to learn about ordering online, but it should not cause them to change preferences for food characteristics.

Data was generously and anonymously shared by a company whose business includes the operation of an online (desktop, mobile, and app) restaurant food ordering business in the U.S. (henceforth, "the platform"). On the platform, consumers find restaurants that operate initially offline, but then start selling additionally via this online channel.[8] Consumers may use homepage product suggestions or engage in active search on the platform to find restaurants in their local area. I study the Los Angeles (LA) metro area market[9] from 2015 to 2018. The platform faced several online competitors during this time period. I will largely abstract from competitive dynamics, though I will discuss how platform competition might affect my estimates below.

Entry onto the platform by existing restaurants comprises much of the variation in assortments in this data.[10] Based on discussions with restaurants and the platform, the typical entrant during this time period was not a new restaurant, but instead was a restaurant already operating offline.[11] New products in this context expand the online assortment but don't also grow the consumer's offline choice set. This will allow me to study how consumers value variety in a specific channel, rather than across all channels or across all related markets.

2.2 Data

I combine several data sources from the platform: (1) two consumer search and purchase panels, (2) restaurant delivery zones, (3) restaurant entry and exit timing, and (4) additional, fixed restaurant characteristics data. The first consumer panel is the set of all new users in the relevant geographic area from January 2015 until June 2016. I use this large, repeated cross-section to study consumer acquisition and immediate retention/churn, as it records the timing and contents of the first order on the platform and how many subsequent orders the consumer made on platform. The second panel is a subset of the initial cross-section, where I follow a 6-month cohort of new users in LA who joined in the first half of 2015. I sample users who order a second time in the first

a restaurant based on consumer choices of menu items. Third, the price can differ because the menu prices change. Finally, the total price can differ when consumers use platform promotional discounts.

[8] Outside of the sample period studied, alternative business models such as "ghost kitchens" have become more widespread, though they are unlikely to affect markets in this paper.

[9] This includes five counties: Los Angeles, Ventura, Orange, Riverside, San Bernardino.

[10] Exit does occur, but it is much less common than entry.

[11] Indeed, the first observed online review for most entrants is months or years prior to their entry onto the platform.

60 days of their lifetime (i.e., don't immediately churn) and remain in LA. I observe their activity in full from 60 days after their first order until September 2018. I sample this second panel more stringently to measure any effects on returning or active customers. The consumer panel is enriched by matching all consumers to Census demographic data at the tract level.

To supplement the consumer order data, I use restaurant characteristics data from the platform and from Yelp. These include cuisine, measures of price and fees, location, entry and exit dates, and matched Yelp data. These data are fixed across time for each restaurant, including measure of price that is not time-varying. Restaurant prices are high-dimensional: each menu item has a price, and the composition of items or the prices attached to them may change or remain constant over time. Item-level pricing is not available. For this reason, I use the average total spending on food and beverage at the restaurant to capture the typical price of buying food at that restaurant.[12] Delivery fees are measured with noise, so I use each restaurant's average fee. To introduce additional price variation, I construct a city-level panel of sales tax rates, which vary over time.

A novel component of the data is delivery zones. During the time period studied, restaurants set their delivery zones in a similar manner to offline or phone-based delivery areas. These zones are relayed in terms of a geographic polygon; a location can receive delivery from the restaurant only if its point contained by the polygon. For the consumer data, I create the realized choice sets at each point in time based on whether the location is in the delivery zone and whether the restaurant is available on the platform.

Table A1 summarizes the 11,286-person returning consumer panel. The median consumer orders 8 times from 4 unique restaurants in the 3.5 years studied. However, there is a substantial right tail of high-consumption users. Figure A6 shows the distribution of choice set sizes for all census tracts where I observe any platform adoption from January 2015 to June 2016. Across the entire geographic area, the median consumer has relatively little choice—the median of this cross-section is 21 restaurants, and the mean is 44 restaurants. These small assortment sizes reflect in part the data construction—census tracts are included even if only a single consumer adopts during the 18 month sample. Such neighborhoods may be in less densely populated, outerlying portions of the LA metro area with low restaurant availability. In contrast, the selection of choice set sizes for the consumer-level panel (shown in Figure A7), a selected sample of consumers, has a median of 85 restaurants. By the end of the panel, this median has grown to 243 restaurants (see Figure A8).

[12] Though prices could change meaningfully over time for some restaurants, data sparsity leads me to average over time.

3 Causal Impact of Assortment Expansion

3.1 Research Design

Unlike many lab studies of assortment sizes and contents, assortments on the platform are not randomly assigned in size or contents. Platform restaurant availability reflects the online and offline decisions of local firms. The set of potential platform restaurants that could serve a neighborhood reflects the equilibrium offline availability of local restaurants. Additionally, these local restaurants choose whether and when to enter the online delivery market. In order to identify the effects of a change in platform assortment on consumers, I will need to address the potential effects of local offline market equilibria and restaurant platform entry strategy.

The number and composition of local offline restaurants reflects the tastes of local consumers, and more restaurants can be sustained by a local market with more consumers and with greater preferences for eating out of home. Restaurants should, all else equal, choose to open in neighborhoods where they expect higher demand. Consumers choose to live in neighborhoods, all else equal, with local amenities that match their tastes (Almagro and Domınguez-Iino, 2021). As a result, consumers who live in neighborhoods with high offline restaurant density are likely to be selected on their preference for restaurant food. That offline density will be reflected in the average number of restaurants available on the platform, since restaurants typically deliver in the immediate vicinity of their physical premises. These factors will then generate positive correlation between the size of the choice set and unobserved consumer or neighborhood heterogeneity in platform behavior.

Platform assortment growth may also coincide with other unobserved demand shocks. Restaurants may time their platform entry around aggregate demand shocks, only entering in periods with high demand. One source of aggregate platform demand shocks is platform promotional activity, which may be timed to coincide with degrees of high assortment growth. Restaurants may select their delivery zone boundaries with potential demand in mind, including areas on the margin where they expect particularly high demand. All three of these forces would generate additional *positive* correlation between assortment growth and platform demand.

My approach to separating this selection from the effect of assortment changes is to condition on neighborhood and time period. To estimate the average effect of an additional restaurant being added to the assortment (entering the platform), I propose aggregating individual restaurant entry quasi-experiments into a staggered entry effect design. I will first present the intuition behind the design, then I will discuss identification conditions.

Consumers (i) living in a neighborhood (z) face a platform choice set of restaurants at time t of S_{it}. As discussed in the prior section, this platform assortments vary across consumers and time, but across-consumer variation is driven only by consumers' order location. My identification approach uses the variation generated by differential entry of restaurants, controlling for consumer or neighborhood unobserved heterogeneity (through fixed effects). I estimate models of the form:

$$y_{it} = \alpha_i + \alpha_{zt} + f(S_{it}, \beta) + \epsilon_{it}$$

$$y_{ct} = \alpha_c + \alpha_{zt} + f(S_{ct}, \beta) + \epsilon_{ct}$$

where c refers to a census tract (level of observation for platform adoption). I primarily study linear effects of size, so $f(S_{it}, \beta) = \beta|S_{it}|$. Outcomes y include adoption, churn, and order frequency.

I consider only within-neighborhood variation in the size of the choice set. Using neighborhood-time fixed effects will isolate variation in restaurant entry to the platform within local areas. The fixed effects will absorb confounding variation generated by restaurant timing selection or any platform promotions.[13] The treatment effect will average across these local comparisons, but will leverage only comparisons between consumers in the same neighborhood who receive different assortment sizes. This variation is comparing consumers who live just on either side of the delivery zone border specified by the entering restaurant.

Identification Example 1. *To better understand the variation that will generate these estimates ($\beta_{adopt}, \beta_{churn}, \beta_{order}$), consider a simple 2-period example. A restaurant ("Z") enters at the beginning of the second period. Half of the consumers, as noted below, are now granted an additional choice on the platform.*

Consumer	*Neighborhood*	*Restaurant Z Availability*
1	*A*	*0*
2	*A*	*0*
3	*B*	*0*
4	*B*	*1*
5	*C*	*1*
6	*C*	*1*

If I did not control for time-varying local markets (with consumer and time fixed effects), the variation used to identify the effect of more choices would comparing all treated users (Consumers 4, 5, 6) to all untreated users (Consumers 1, 2, 3). If orders are the outcome of interest, such an

[13]During this time, promotions did not follow detailed geographic targeting.

approach would compare changes in orders between treated and untreated users. In the specification I will present throughout with neighborhood-time fixed effects, the regression will only use variation in entry that varies within a neighborhood-time period. In this example, only the variation in choice sets in Neighborhood B will be used, since the neighborhood-time fixed effects will soak up variation in Neighborhood A and Neighborhood C.

I want to highlight the residual variation used in this context under this strategy. The two-way fixed effects approach soaks up nearly all of the variation in choice set size in this data. Table OA17 shows the R-squared from regressions of different two-way fixed effect regressions on the treatments of interest: the assortment size in levels and changes to the assortment size. In levels, the fixed effects explain nearly all of the differences across consumers and time in the size of the choice set. However, this is slightly misleading. The marginal effect of the assortment size is identified here from changes in assortments. The fixed effects explain a considerable share of the changes to assortments: over 90 percent of the entry is explained by neighborhood-week fixed effects alone. Despite the size of the data, I should expect results to be relatively low powered given the share of actual variation used to identify this main effect. The set of residual variation used is quite small, so I will address endogeneity concerns specifically with the residual, identifying variation in mind. Because I leverage assortment variation determined at the local neighborhood market - time level, I cluster standard errors at the neighborhood market-time level based on varying propensities for treatment (Abadie et al., 2017).

3.2 Identification Assumptions

Given this design, I assume that the assortment size is independent of the unobserved determinants of consumer behavior, conditional on consumer and time specific heterogeneity. Specifically, $S_{it} \perp \epsilon_{it} | \alpha_i, \alpha_t$. Conditional on consumers' locations, their persistent taste heterogeneity, and any aggregate time effects, variation in the size of the platform assortment is assumed to be independent of the error term(s). Implicit in this assumption is that, absent changes to the assortment, consumers' behavior on the platform follows parallel trends.

Conditional on consumer geographic selection, the main challenge to identification is strategic behavior on the part of restaurants. In this empirical context, restaurants cannot precisely control their exact entry timing, so they cannot select entry timing to coincide with positive weekly demand shocks. However, the shape and size of their choice of delivery zones could potentially violate the identification assumption. The overall size of these zones is fairly uniform, with typical radii around the physical restaurant location of 3 to 4 kilometers. Strategic behavior in drawing the exact boundary, conditional on approximate size, could create problems for identification. Restaurants may choose to select their boundaries by including blocks where they expect to sell and excluding blocks with low demand, on the margin. If this is the case, then the estimates here will be an upper bound on the true effect, since such strategic behavior would generate positive selection.

These are strong assumptions and worth discussing in practical detail. In particular, I want to emphasize what strategic behavior by restaurants and the platform is *ruled out* by this design. I will additionally discuss assumptions about dynamic treatment effects that are testable and ruled out. These assumptions would be violated by many forms of platform personalization or geographic based targeting. I am assuming that the platform does not target promotions geographically in a way that is related to entry, and that restaurant entry does not change the platform design except via the size of the assortment. During this time period, the platform studied did not engage in such fine-grained targeting.

I am implicitly assuming that platform promotional activity (advertising or discounts) only varies across consumers independent of changes to the assortment. In particular, this rules out that the platform engages in targeting on past treatment effects. For example, if β is positive, I assume that the platform does not send promotions to remedially improve adoption or retention in areas with low assortment size. If β is negative, I assume that the platform doesn't remedially target areas or consumers with high assortment growth with promotions. However, conditional on past entry (and any consumer dynamics), I am ruling out that *entry itself* alters platform promotions, which is consistent with practice at the time.

One additional source of unobserved, systematic variation in platform adoption and consumption is platform competition. The platform studied here has a handful of similarly sized competitors during the sample period. Any promotional activity by competing platforms—so long as it is not geographically targeted within a metro area—will be soaked up by time fixed effects. As noted above, geographic targeting with this level of precision was not common practice during this period. The other way platform competition could create a problem for this identification strategy is if restaurants simultaneously enter multiple platforms with identical delivery zones. If restaurants were to enter in this manner, it would be impossible to attribute the effect observed on this platform solely to platform assortment changes. However, based on conversations with the platform, such simultaneous entry is unlikely. I measure entry across platforms based on consumer review text (detailed in Appendix OA5), and I reject that restaurants are entering this platform's two largest competitors around the same time.

Finally, the difference-in-differences strategy rules out dynamic treatment effects. However, some of these dynamic effects can be included by testing estimating equations that include treatment lags or cumulative measures of past treatment changes. In light of recent work highlighting potential pitfalls of two-way fixed effects for estimating difference-in-differences research design (Sun and Abraham, 2020; Callaway and Sant'Anna, 2020; De Chaisemartin and d'Haultfoeuille, 2020), I also test robustness to alternative estimators in Appendix OA2.

3.3 Results

Larger platform assortments increase the number of adopting (first-time) consumers but reduce the frequency with which returning consumers order on the platform. Conditional on adoption, the size of the assortment at the point of adoption does not significantly alter the probability of churn after the first order. Table 1 summarizes the direction of the estimated effects, and Table 2 shows estimates for the main specifications for the three outcome measures mentioned and results which control for the degree of assortment variety.

Table 1: Summary of Effects

	Sign	Effect
$\frac{\partial A_{ct}}{\partial \lvert S_{ct} \rvert}$	(+)	Large assortments increase adoption
$\frac{\partial C_{ct}}{\partial \lvert S_{ct} \rvert}$	(0)	Larger assortments at adoption don't impact immediate churn
$\frac{\partial o_{it}}{\partial \lvert S_{it} \rvert}$	(-)	Larger assortments reduce order frequency among returning users
$\frac{\partial s_{it}}{\partial \lvert S_{it} \rvert}$	(-)	Larger assortments reduce search sessions
$\frac{\partial o_{it}}{\partial \lvert S_{it} \rvert} \lvert s_{it} > 0$	(0)	Larger assortments don't impact orders conditional on search

Column 1 of Table 2 shows the estimated marginal impact on an additional restaurant on the platform on the platform's adoption rate at the census tract level. The addition of 10 more restaurants is estimated to increase the adoption rate by 0.002 percentage points, which is a 5-7 percent increase over the baseline. This effect is consistent once controlling for the assortment's variety (measured as the number of unique cuisines), as reported in Columns 2 of Table 2. The small, positive indirect network effect found here—as measured by the total number of users signed up for the platform—is consistent with past work that documents small, positive effects in other contexts.

These effects on adoption (the installed user base) are potentially driven by diffusion from existing, larger seller (restaurant) bases and by the entry of new restaurants. Columns 1 and

Table 2: Main Effects of Assortment Size

	Dependent variable:					
	Adoption Rate		First-Time Churn Rate		Orders	
	(1)	(2)	(3)	(4)	(5)	(6)
Restaurant Count (tract)	0.0002^{**}	0.0002^{**}	-0.023	-0.020		
	(0.0001)	(0.0001)	(0.030)	(0.030)		
Cuisine Count (tract)		-0.0002^{*}		0.347^{*}		
		(0.0001)		(0.175)		
Restaurant Count (household)					-0.0001	-0.0001^{**}
					(0.00003)	(0.00003)
Cuisine Count (household)						0.002^{***}
						(0.0004)
Elasticity of # Rest.	0.3182	0.3235	-0.0579	-0.0498	-0.1146	-0.1920
Elasticity of # Cuis.		-0.189		0.359		0.483
Observations	204,798	204,798	104,218	104,218	2,058,406	2,058,406
R^2	0.776	0.776	0.331	0.331	0.237	0.237
Adjusted R^2	0.740	0.740	0.114	0.114	0.207	0.207

Note: $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$
All specs include ZCTA-Week and Unit FEs.
Standard Errors are clustered at the ZCTA-Week Level.
Churn Rate results omit promotion use control.
Rates range from 0 to 100

Table 3: Effects of Entry and Lagged Assortment Size

	Dependent variable:					
	Adoption Rate		First-Time Churn Rate		Orders	
	(1)	(2)	(3)	(4)	(5)	(6)
Restaurant Entry (tract)	−0.0002	−0.0002	0.033	0.088		
	(0.0003)	(0.0003)	(0.258)	(0.264)		
Restaurant Entry (household)					-0.005^{***}	-0.005^{***}
					(0.0004)	(0.0004)
Lag Restaurant Count	0.0002^{**}	0.0002^{**}	−0.032	−0.028	−0.00003	-0.0001^{*}
	(0.0001)	(0.0001)	(0.031)	(0.031)	(0.00003)	(0.00003)
Cuisine Entry (tract)		0.0003		−0.463		
		(0.0003)		(0.617)		
Cuisine Entry (household)						−0.001
						(0.001)
Lag Cuisine Count		-0.0002^{*}		0.331		0.002^{***}
		(0.0001)		(0.179)		(0.0004)
Observations	201,695	201,695	102,928	102,928	2,058,406	2,058,406
R^2	0.778	0.778	0.331	0.331	0.237	0.237
Adjusted R^2	0.743	0.743	0.115	0.115	0.207	0.207

Note: $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$
All specs include ZCTA-Week and Unit FEs.
Standard Errors are clustered at the ZCTA-Week Level.
Churn Rate results omit promotion use control.
Rates range from 0 to 100

2 Table 3 decomposes the average effect of assortment size. The existing assortment size seems to drive much of the effect—estimates for the impact of restaurant entry are mixed. However, new users are twice as likely to have ordered from a recent (within 30 days) entrant than more experienced users. Specific entering restaurants may drive their existing offline customers to join the platform, and more extant variety also drives adoption. In contrast to other outcome measures, the scope of variety (measured by number of cuisines) does not enhance adoption.

The size of the assortment at the moment of adoption may impact the retention of new customers *after* adoption. To measure this, I estimate the effect of assortment size at adoption on whether the consumers *immediately* churn from the platform after the first order. These estimates are shown in columns 3 and 4 of Table 2 and in columns 3 and 4 of Table 3 with additional controls. The estimates are mixed in sign across specifications. These results omit the coefficients from a control: the average share of the initial order purchased under promotion. The estimated coefficients on this promotion usage variable are large, positive and significant—consumers who use coupons when adopting are considerably more likely to churn after their first order.

The estimated effect on returning customers' order frequency is presented in columns 5 and 6 of Table 2. For returning customers, adding 10 restaurants to the platform in their area reduces their weekly orders by 0.001, or about 1 percent. In the main specification, the estimates are not statistically significant.[14] However, once I control for a measure of product variety—the number of cuisines offered—this effect is significant. Though the net effect of variety is small in magnitude, it may reflect larger costs and benefits which act in opposing directions. In fact, larger negative effects would be potentially implausible given consumers could ignore entry and consume only from incumbent sellers. Indeed, actual expansion of variety by offering new product attributes increases consumption, and such entry more than offsets the negative crowding effect measured in the number of restaurants.

The source of variation in assortments on the platform is restaurant entry[15] onto the platform. Columns 5 and 6 of Table 3 compares the main specification (effect of total number of restaurants on orders) to an additional specification which separates the effect into entry (or exit) and lagged assortment size. The entry of new restaurants significantly reduces weekly orders—one new restaurant entering reduces the number of orders by 0.005, or 5 percent. The size of the existing assortment only has a small negative effect.

[14] I can rule out a positive effect of assortment size on consumption frequency as predicted by typical demand models: the one-sided test (rejecting a null hypothesis of $\beta_{order} > 0$) is statistically significant.

[15] Exit also occurs, but is relatively infrequent. Separate estimates of exit and entry effects are similar in magnitude.

The platform may care about order frequency, but its revenue is also dependent on basket size. The entry of additional restaurants reduces the average spending, as shown in Table OA6. This spending reduction primarily occurs due to reduced frequency—the average size of orders, conditional on ordering, is slightly higher as the assortment grows.

Nonlinear effects across the choice set size are limited. Table OA7 presents regression results that allow the marginal effect of an additional platform restaurant to differ across five assortment size bins. These results are consistent with the uniform effect—the marginal negative effect of additional restaurants is similar across choice set sizes.

The reduction in consumption occurs by reducing experimentation among returning consumers. Assortment growth significantly reduces ordering from never-before-tried restaurants, but not from already-tried restaurants (see Table 4).

Table 4: Effect of Assortment Size on Orders by Type

	Dependent variable:	
	First-Time Orders	Repeat Orders
	(1)	(2)
Restaurant Count	−0.0001***	0.00001
	(0.00002)	(0.00002)
Observations	2,058,406	2,058,406
R^2	0.100	0.228
Adjusted R^2	0.064	0.197
Note:		*p<0.05; **p<0.01; ***p<0.001
		All specs include Individual and ZCTA-Week FEs
		Standard Errors clustered at ZCTA-week level

Consumers may also purchase less frequently if large choice sets lead them to buy lower quality or more expensive products due to the difficulty of finding products. Table 5 shows how the average characteristics of restaurants ultimately chosen vary with the size of the assortment. When customers do ultimately purchase from a larger assortment, they chose restaurants which are, on average, less popular overall, as measured by the total number of Yelp reviews, but may be of similar popularity on the platform.[16] There is not a consistent effect across specifications on the quality of the restaurant (measured as 4.5 or 5 stars on Yelp) or the average price of a basket of food at the restaurant. The selection of restaurants that are less popular offline may not necessarily

[16]This popularity measure is constructed as the sales quantile of the restaurant among this cohort from the entire panel. It does not condition on availability.

reflect lower quality—successful offline restaurants may be popular due to the quality of in-person service, which does not translate to the quality of delivery service.

Table 5: Effect of Assortment Size on Ordered Restaurant Characteristics

	Characteristics of Ordered Restaurant			
	Platform Popularity	Yelp # Reviews	Yelp Rating > 4	Avg. Price
	(1)	(2)	(3)	(4)
Restaurant Count	−0.0001	−0.653**	−0.0001	0.003
	(0.0001)	(0.288)	(0.0001)	(0.003)
Observations	121,820	121,820	121,820	121,820
R^2	0.628	0.555	0.555	0.616
Adjusted R^2	0.408	0.293	0.292	0.390
Note:				*p<0.1; **p<0.05; ***p<0.01
				All specs include ZCTA-Week and Individual FEs
				Standard Errors clustered at ZCTA-Week level

All restaurant entry may not impact consumers equally. I consider three types of heterogeneous effects of different types of restaurant entry on order frequency: restaurant chain status, restaurant vertical quality rating, and restaurant match with individual consumers. Consumers may have very different information or consideration costs from chain restaurants than local independent restaurants. In particular, I expect that there consumers would have little need to search over national or prominent local chains. Consistent with this explanation, the number of independent restaurants significantly reduces order frequency, but the number of chain restaurants has minimal or positive effect (shown in Table OA1).[17]

If new restaurants are lower-quality than incumbent restaurants, consumers expectations of the value of ordering on the platform may be diluted, driving lower return frequency. Given the reduction in purchases from novel alternatives (Table 4), I expect that adding low-quality entrants would reduce consumption frequency more than high-quality entrants. I test this by breaking up the assortment by Yelp star ratings. Table 6 presents the estimated effect of assortment size by binned Yelp ratings. The results are noisy but suggest this effect is not ameliorated by high-rated restaurants entering the assortment. In particular, adding very-highly rated restaurants to the platform still reduces consumption frequency. These results are imprecise—I cannot rule out a

[17] Chain restaurants include large, national quick-serve and fast casual restaurants, regional chains, and local chains with at least 5 outlets.

small positive effect from low quality entrants on consumption frequency. I also cannot rule out that restaurants of all vertical quality ratings reduce the probability of consumption on the platform by returning users.

Table 6: Effect of Assortment Quality on Returning Customers

	Dependent variable:	
	Weekly Orders	Weekly Spending (USD)
	(1)	(2)
Restaurant Count (4.5 or 5 stars)	-0.001^{***}	-0.009
	(0.0002)	(0.005)
Restaurant Count (3.5 or 4 stars)	0.0002^{*}	0.006^{**}
	(0.0001)	(0.002)
Restaurant Count (3 stars or less)	-0.0003	-0.009
	(0.0002)	(0.006)
Observations	1,991,134	1,991,134
R^2	0.239	0.222
Adjusted R^2	0.209	0.191
Note:		*p<0.05; **p<0.01; ***p<0.001
		All specs include ZCTA-Week and Individual FEs
		Standard Errors clustered at ZCTA-week level

Even if there is no effect of vertical quality, entry could still dilute individual consumer expectations about match value if the changes to the assortment are mostly low-match-value products for their particular tastes. To proxy for this, I distinguish between relevant (ever consumed) and irrelevant (never consumed) cuisines for each consumer. This proxy may be noisy. For a consumer who orders pizza, the addition of more pizza restaurants may be irrelevant, as they already have found a preferred pizza restaurant. Conversely, a consumer who never orders pizza on the platform may still consider it for purchase. Table 7 shows the marginal effect of relevant- (category consumed) versus irrelevant- (category never consumed) restaurants added into the assortment. I find that the negative effect on purchase frequency is driven by growth in relevant restaurants. The addition of restaurants which are less relevant, in contrast, increases the probability of purchase.[18]

Poor experiences purchasing from a large assortment could potentially drive consumers to churn from the platform, though they remain observed in the panel. To consider whether this is driving my results, I look at three additional analyses. First, I control for individual-year fixed effects, and

[18]This positive effect could be driven by positive externalities from entry of substitutes which are irrelevant to the consumer. Entry in other cuisines could *reduce* crowding in a consumer's preferred restaurants.

Table 7: Impact of Relevant Restaurant Entry on Returning Customers

	Dependent variable: Weekly Orders	
	(1)	(2)
Ever Consumed Cuisine Restaurants	−0.001***	−0.001***
	(0.00003)	(0.00004)
Never Consumed Cuisine Restaurants		0.0003***
		(0.00003)
Observations	2,058,406	2,058,406
R^2	0.239	0.239
Adjusted R^2	0.209	0.209
Note:		*p<0.05; **p<0.01; ***p<0.001
		All specs include Individual and ZCTA-Week FEs
		Standard Errors clustered at ZCTA-week level

I find consistent results. As consumers churn, the lack of orders will be fully absorbed by these fixed effects. Second, I look at the effect of assortment size for consumers who still make an order in the final year of data, and I find that assortment size decreases order frequency. Finally, I look descriptively at when consumers churn from the panel. Most churn occurs in the first year panel. Given this is the case, the results using more granular user-time fixed effects should control for any exodus of users.

Consumer Search in Larger Assortments

To disentangle the negative effect of variety on returning customers further, I supplement the order data with limited summary data on search behavior for a subset of consumers.[19] Does the growth of assortments lead consumers to search more? Does longer search lead consumers to learn about the cost of finding a good option in large choice sets?

I observe weekly counts of search sessions on the platform, which allows me to construct conversion rates conditional on search. Using this subset of about half of the consumers, I document that the elasticity of searching with respect to assortment size is about -0.5: the addition of 1% more restaurants reduces weekly search by 0.5%. Conditional on searching, however, there appears to be no or a small positive impact of assortment size on search conversion into ordering. Table 8 shows the effect of assortment size on weekly search sessions and total search duration. This reduction at the 'top of the funnel' is inconsistent with the in-person choice overload experiments, where greater

[19] Appendix OA4 details selection of users into the search data.

in-person assortments draw consumers' attention at higher rates (Iyengar and Lepper, 2000).

Table 8: Effect of Assortment Size on Consumer Search Behavior

	Dependent variable:	
	Weekly Search Sessions	Weekly Search Duration
	(1)	(2)
Restaurant Count	-0.0002^{**}	-0.002
	(0.0001)	(0.001)
Observations	1,436,956	1,436,956
R^2	0.262	0.170
Adjusted R^2	0.222	0.126
Note:		*p<0.05; **p<0.01; ***p<0.001
		All specs include ZCTA-Week and Individual FEs
		Standard Errors clustered at ZCTA-week level
		Search duration measured in minutes

Table 9: Lagged and Contemporaneous Effects of Assortment Size on Consumer Search Behavior

	Dependent variable:	
	Weekly Search Sessions	Weekly Search Duration
	(1)	(2)
Restaurant Entry	-0.011^{***}	-0.077^{***}
	(0.001)	(0.015)
Lag Restaurant Count	-0.0001^{*}	-0.001
	(0.0001)	(0.001)
Observations	1,436,956	1,436,956
R^2	0.262	0.170
Adjusted R^2	0.222	0.126
Note:		*p<0.05; **p<0.01; ***p<0.001
		All specs include ZCTA-Week and Individual FEs
		Standard Errors clustered at ZCTA-week level
		Search duration measured in minutes

This decrease in search differs by what ultimately results from the search session. Consumers who search and ultimately purchase spend longer searching pre-purchase in a larger assortment, but those who search and then fail to convert have shorter sessions as the assortment grows (Table 10), though these results condition on having searched (which occurs less frequently in due to assortment growth). However, they are consistent with some accounts of choice overload—where consumers avoid search if overwhelmed. Moreover, consumers search longer in larger assortments when they ultimately choose an untried restaurant, perhaps reflecting the ease of accessing past choices on the platform (Table OA3).

Table 10: Effect of Assortment Size on Search Duration by Purchase Status

	Dependent variable:	
	Search Duration (Minutes)	
	(1)	(2)
Restaurant Count × No Purchase	−0.015	−0.016
	(0.022)	(0.022)
Restaurant Count × Purchase	0.047*	0.043*
	(0.022)	(0.022)
Selection Controls?	N	Y
Observations	84,630	84,630
R^2	0.556	0.569
Adjusted R^2	0.226	0.248
Note:	*p<0.05; **p<0.01; ***p<0.001	
	All specs include ZCTA-Week and Individual FEs	
	Omits search selection first stage residual control	
	Standard Errors clustered at ZCTA-week level	

3.4 Robustness Checks

I conduct three main robustness checks to verify whether these causal effect estimates are robust to alternative explanations: strategic restaurant entry timing, limited updating by consumers, and variation in entrant quality.

Merger Natural Experiment

I use a platform merger as a natural experiment to check that the effects are robust to otherwise endogenous entry timing. The platform studied in this paper, during the sample, acquired several smaller competing platforms. After the merger was completed, the platform on-boarded the restaurants from the acquired platforms and released them online on a handful days during this time period. These discrete jumps serve as a natural experiment, since the entry timing of these restaurants was not based on restaurant strategic behavior or on platform strategic behavior. The results from these natural experiments will be less precise, since the individual fixed effects cannot be estimated with precision in short panels, and less than 3% of the original panel is used for estimation. The results of this robustness check are partially consistent with previous findings, but they are noisy null results (Tables OA12-OA14). The effect of entry on returning customers' order frequency using only these few observations is positive.[20]

Alternative Assortment Measures

I check the robustness of the estimates to consumer limited information by considering the extreme where consumers only update their understanding of assortment size when they interact with the platform. I construct the alternative size of the assortment at the last search time period, and I carry it forward. Additionally, given the average growth trajectory on the platform, I allow an approximate linear updating of the assortment over time in line with this growth.

In both cases, I also control for the time since last search. Simply repeating the main specifications with this new measure will generate selection that will bias the estimates. The platform is generally growing over time, so households who more recently went online are going to have higher assortment sizes, all else equal. However, their recent purchase or search also can reflect higher engagement and purchase likelihood overall, which could lead to spurious positive correlation between the size of the assortment and purchase probability.

Results from these alternative measures (Tables OA15), which would allow for the possibility that consumers aren't fully aware of assortment changes, are consistent with a marginal negative effect on weekly ordering. This interpretation, however, does not square exactly with the prior finding that concurrent assortment changes negatively impact ordering. Because growth of the assortment is correlated across time within an area—i.e. high growth areas remain so throughout the sample —it's not possible to directly test the 'lagged perception' against current changes. I estimate a version of the specification with both measures. The effect of the concurrent assortment size on consumption remains significant and negative, while the alternative measure now has a noisy null effect. From these, I conclude that the marginal effect of assortment growth may have some spillovers over time, and the effect is robust to alternative consumer updating frequency.

Differential Effects across Restaurants

Based on the findings in Ershov (2022), I consider that new additions to the choice set may be meaningfully worse quality than incumbents. If new restaurants are worse, the average product quality could decline, which could partly explain the decline in consumption I have documented—though this could be soaked up by time fixed effects. This mechanism should be more muted in my context, as the introduction of low-quality new products does not, in principle, affect the consumption value of existing preferred products. Using observable quality measures, I found above

[20]Though this effect is inconsistent in sign with other estimates, the confidence interval contains point estimates from the full sample and it is statistically indistinguishable from zero or a negative effect.

that a high-quality new entrant is equally detrimental as a low-quality new entrant.[21]

I observe some measurable differences in attributes when comparing entrants (restaurants that enter during the sample) and incumbents. Table OA16 shows the distribution of product attributes among incumbents versus entrant restaurants. The prices charged by these restaurants differ slightly, but this occurs only in terms of delivery fee versus food costs. The total cost is very similar across the two groups. Older (i.e. incumbent) restaurants have many more reviews on average than new restaurants, though this is unsurprising as they have had longer to accumulate them. Incumbents are marginally higher rated on average than entrants. This may reflect selection on surviving restaurants: the incumbents that remain into the panel are ones that have not yet closed. New restaurant entrants also have lower sales, on average, than existing restaurants on the platform. This quality selection could contribute to quality dilution by new entrants, which in turn may contribute to the negative impact of assortment expansion on purchasing.

I conclude that the negative effect is not driven by higher prices or lower observed quality from new entrants, and that it is possible some of the effect is driven by unobserved restaurant quality. However, the marginal negative effect on overall orders in the repeated-purchase context cannot be fully explained by quality in a standard demand model. In particular, even if the average entrant is of lower (unobserved) quality, the negative effect on sales of *existing incumbent* restaurants is a violation of the independence of irrelevant alternatives (IIA).

3.5 Discussion

In this section, I showed how larger assortments marginally improve customer acquisition and reduce consumption among existing customers. Though results are consistent in direction and magnitude across a variety of specifications, in many cases the results are not statistically significant. In the case of such small effect sizes this is not surprising, but I want to emphasize what can be concluded from these imprecise estimates. I can rule out large positive effects of product entry for returning customers, though it is possible that the true effect of entry is zero.

Reduced consumption occurs through increased interpurchase time, not smaller baskets. Descriptive evidence is consistent with the presence of consumer search frictions and incomplete information. Search duration (the total time cost expended prior to purchase) is higher when choice sets are larger. This, in turn, increases the time between purchases (Table OA2). Consumers search longer prior to a given purchase, which then increases the time until they subsequently return to the

[21] Yelp rating may be a very noisy measure of quality. Even if that is the case, this information is displayed to consumers on this and many other platforms.

platform. This effect is strongest following consumers' purchase from a novel-to-them restaurant (Table OA3). Many of these results are statistically insignificant, but I can rule out large effects in the opposite direction.

These results are inconsistent with full information demand models, but the exact mechanism by which information costs are higher under larger assortments is unclear. One possibility, as in Kuksov and Villas-Boas (2010), is that consumers update their expectations about total information costs, though the per-product information costs are unaltered. This would also be consistent with consumers' expectations about match value changing with the size of the assortment. Another possibility is that per-product information costs are higher, since consumers have to sort through more products to acquire information about any particular restaurant.

I rule out several mechanisms through the research design and through robustness checks. First, this reaction is not through observable quality or price differences (Ershov, 2022) between entrants and existing restaurants. Second, since I have granular time and area-time fixed effects, I rule out that these effects stem from platform-level promotional activity, which did not have any detailed geographic targeting component during this time period. Third, across all users, conditional on ordering, larger assortments induce higher rates of repurchase/lower rates of experimentation. Given the structure of the user experience, this is consistent with users potentially gathering information in a less costly manner, by navigating from the home page which often presents recently ordered-from restaurants.

Though the body of evidence from these reduced form approaches rules out positive effects of entry for returning consumers, the average effects measured cannot distinguish between minimal effect of variety and two counteracting forces. Choice overload could be a very small effect in concert with no benefit from variety. However, the crowding effect could be much larger, but the regressions only measure the combined crowding effect and the benefit of variety. To think about the efficacy of alternative assortment strategies, I need to decompose these two effects. In the next section, I will build a structural model of consumer demand with two aims. First, I will distinguish between multiple ways in which the search process could be altered by changes to the assortment. Second, it will measure heterogeneous preferences, so that consumers' choice of particular restaurants can related to product characteristics. This will allow me to consider how removing any particular restaurant from the choice set will impact an individual consumer's purchase behavior.

While there are many benefits to the specific empirical context, the structure of the consumer purchasing decision (discrete choice) does prevent me from capturing the full benefit of variety.

Grocery retail, where consumers typically purchase many goods that comprise a basket across categories, allows for a better measurement of the returns to consuming multiple goods. My interpretation of these results in broader contexts is that these negative effects of large assortments may be harder to detect, but still influence consumer behavior. The other limitation of this context is that I use only a narrow cohort of users, and there may be some adoption-time specific effects.

4 Limited Information Demand and Counterfactuals

Larger assortments reduce the purchase frequency of returning customers, but the prior section does not provide a clear remedy for platforms or retailers. I build a structural model of consumer information acquisition and demand which nests a test of the mechanism by which larger assortments reduce purchasing. Distinguishing this mechanism (along with estimating consumer preferences) is necessary in order for platforms and retailers to address the reduced purchasing among existing customers. Reducing assortments by removing high match products could potentially reduce sales by much more than any congestion or choice overload effects. Without separating product preferences from congestion effects, platforms are unable to engage in effective assortment reductions. I distinguish between assortments altering the cost of learning product information from altering consumers' expectations about product match. These mechanisms, while both directly addressable by reducing the assortment size, suggest different paths for how else platforms might improve retention. I use the model results to test several assortment reduction strategies. Offering smaller assortments to each consumer improves the expected revenue to the platform only if the reductions are targeted based on consumers' preferences.

4.1 A Model of Discrete Choice under Incomplete Information

A consumer i faces a choice set of restaurants S_{it} on the platform at time t based on their location. After adoption, consumers choose to order from a single restaurant $j \in S_{it}$ (measured by choice dummy y_{ijt}) or the outside option (denoted $j = 0$) every period. Consumers have heterogeneous tastes over restaurant attributes. Their consumption utility from each option is $u_{ijt} = \delta_{ijt} + \zeta_{ijt}$. I assume consumers have incomplete information about each products consumption utility, though they know the contents of their choice set (and thus its size). Consumers know some product information δ_{ijt} costlessly, but not ζ_{ijt}. ζ is mean-zero across alternatives by construction.

I assume consumers are rationally inattentive (e.g. Sims, 2003; Matêjka and McKay, 2015) and acquire additional information, which is costly to them, prior to choosing. My model (and its exposition) follows closely Joo (2022), where consumers have subjective expectations about product valuations and product attributes are non-stochastic.

Each period, consumers arrive and are endowed with an information structure (collection of vectors $\mathbf{D_{it}} := (\mathbf{D_{i1t}}, \ldots, \mathbf{D_{iJt}}))$, including the size and contents of the assortment and the information cost function. Upon arrival, consumers form subjective prior beliefs over consumption utilities of each restaurant based on the endowed information $\mathbf{D_{it}}$ such as their own purchase history and promotional activity (e.g., prominence on the platform). These beliefs are allowed to be subjective since, unlike Matêjka and McKay (2015), consumption utility (u_{ijt}) may be deterministic and fixed over long periods of time. Consumers are instead uncertain about which products map to which utilities. Denote these subjective prior distributions over the vector of consumption utilities $Q_i(\cdot) = Q(\cdot||S_{it}|, \mathbf{D_{it}})$. Prior to attending to product information, consumers have ex ante unconditional probabilities of choice for each alternative

$$\pi_{ijt} = \int Pr(i \text{ chooses } j \text{ in } t|\mathbf{u})dQ(\mathbf{u}||S_{it}|, \mathbf{D_{it}})$$

Given their prior beliefs, consumers design an information collection strategy, which weighs the expected improvement in choice quality against the cost of attention. The intuition of this optimization step is very similar to fixed-sample search strategies (e.g. Chade and Smith, 2006), since the consumer chooses the attention strategy upfront and does not alter it sequentially in response to information gained. After attending to alternatives according to their information collection strategy, consumers update their beliefs in a Bayesian manner and chooses the alternative that maximizes their payoff based on posterior beliefs. Formally, the attention strategy is the solution to

$$\begin{aligned} \max_{\{Pr(i \text{ chooses } j \text{ in } t|\cdot)\}_{j\in S_{it}}} \int \{ \sum_{j\in S_{it}} u_{ijt} Pr(i \text{ chooses } j \text{ in } t|\mathbf{u}) \\ -c(\pi\{Pr(i \text{ chooses } j \text{ in } t|\mathbf{u})\}_{j\in S_{it}})\} dQ(\mathbf{u}||S_{it}|, \mathbf{D_{it}}) \end{aligned}$$

where the choice probabilities sum to 1 and include the outside option. This strategy is the solution in equation 4 of Joo (2022). The first term is the expected benefit from choice, and the second term is the information cost function. This function is assumed to be proportional to the reduction in (Shannon) entropy in consumers' beliefs between the prior belief and the post-attention posterior belief.[22]

$$c(\pi\{Pr(i \text{ chooses } j \text{ in } t|\cdot)\}_{j\in S_{it}})$$
$$= \frac{1}{\mu}\left[\sum_{j\in S_{it}} Pr(i \text{ chooses } j \text{ in } t|\cdot)ln(Pr(i \text{ chooses } j \text{ in } t|\cdot)) - \sum_{jinS_{it}} Pr(i \text{ chooses } j \text{ in } t|\cdot)ln(\pi_{ijt})\right]$$

where the unit cost of information has inverse μ_{it}. This function, with roots in information theory, is convex—it is increasingly costly for consumers to gather additional information that would reduce uncertainty across choices. While it is possible for a consumer to reach full information in this model, any nonzero information costs would make this degree of attention extremely costly relative to improvement in choices.

Given the prior beliefs, the choice of attention strategy, the information cost function, and the post-attention choice rule, Joo (2022) shows that consumers' choice probabilities are

$$Pr(i \text{ chooses } j \text{ in } t|\mathbf{u_{it}}) \equiv P_{ijt}(\mathbf{u}_{it}) = \frac{exp(ln(\pi_{jt}) + \mu_{it}u_{ijt})}{1 + \sum_{k\in S_{it}} exp(ln(\pi_k t) + \mu_{it}u_{ikt})} \tag{1}$$

Given the delineation of utility into known and unknown components in my model, I can write more specifically the probability that consumer i chooses restaurant j in period t as:

$$P_{ijt}(\mathbf{u}_{it}) = \frac{\pi(\delta_{ijt})exp(\mu_{it}u_{ijt})}{1 + \sum_{k\in S_{it}} \pi(\delta_{ikt})exp(\mu_{it}u_{ikt})} \tag{2}$$

I select this framework for consumer information acquisition (as opposed to other models of consumer search) for its flexibility and tractability in large choice sets. Unlike typical sequential search models, attention does not result in all-or-nothing product information—greater attention to an alternative provides more information.

4.1.1 Impact of Assortment Expansion

The basic rational inattention framework for discrete choice does not necessarily encompass the effects of asssortment size documented above. Adding a new product to the choice set, as shown in Joo (2022), increases the probability of purchase but may decrease consumer welfare. I need the model to be even more flexible to allow for the possibility of reduced inside choice shares. In

[22] As noted in Joo (2022), under the subjective prior RI model, the interpretation of information costs differs slightly from the rest of the RI theory literature. In particular, it should be thought of as"the cost of consumers' choice adjustments associated with changing the choice probabilities from unconditional choice probabilities... to conditional choice probabilities" (Joo (2022), page 40). In this context, one can interpret this as the combined cost of attending to information and adjusting planned decisions.

particular, while adding products to the assortment may add new characteristics, they also may alter consumers' unconditional probability of purchase ($\pi(\delta)$) via expectations about choice-specific match value. Additionally, larger assortments may alter the cost of accessing product information. In either case, this may alter whether consumers make any purchase and which products they choose to purchase. I assume that the assortment does not directly alter consumption utility of products, conditional on choice.[23]

These two channels (expectations and information costs) by which assortment size impacts consumers' information acquisition and product choice have distinct predictions for consumer behavior. If the cost of searching products increases with the size of the choice set, ceteris paribus, consumers will become less sensitive to the hidden portion of product utility. This could be consistent with consumers making higher-price or lower-quality selections, conditional on purchase. The impact of assortment size on consumers' expectations is less straightforward. If the restaurant equilibrium was modelled, we could recover how average match value might change with the size of the assortment. Absent that, however, assortment growth that alters expectations does not change how sensitive consumers are to the post-search characteristics, ceteris paribus.

These channels by which the contextual information about the size of the choice set impacts choice—prior beliefs about value, and information costs—reflect existing theoretical explanations for choice overload effects. Kamenica (2008), Kuksov and Villas-Boas (2010), and Nocke and Rey (2022) provide accounts for how larger numbers of products reduce the probability of choice. In Kamenica (2008) and Kuksov and Villas-Boas (2010), a larger number of products implies that the consumer will be worse off in expectation from consumption, either because the match value of the product is worse, or because the expected information costs outweigh the benefits from a better match value. Nocke and Rey (2022) find that, in equilibrium, firms have incentives to reduce the informativeness of product orderings conditional on assortment size (garbling overload), and increased assortment size discourages consumers from engaging with the product line because of this information garbling.

The average information cost of each product also could be altered by the total number of products due to more total search results from querying or due to higher information processing costs in the presence of more products. First, consumers have to sift through more search results, on average, to reach any product; search costs increase considerably with search result position

[23] This contrasts with the typical use of congestion terms in indirect utility functions, in the style of Draganska and Jain (2005).

(Ursu, 2018). Second, the cost of searching more intensively (paying more attention cost) may be altered by the amount of information displayed in total (Chandon et al., 2009; Gu and Wang, 2022).

Previous empirical implementations of rational inattention do not accommodate my descriptive findings—that the probability of purchasing any product declines for returning consumers in the size of the choice set. Given the assumption on information costs, the logit-like choice probabilities suffer from the same problem as a standard multinomial logit. Holding price and other attributes fixed, the addition of an alternative weakly improves the probability of selecting an inside option. Unless the size of the assortment directly enters into either $\pi_{i,k}$, μ_i, or $u_{i,k}$, the size of the assortment can only increase the probability of purchase.

To address this, I assume that the cost of information and beliefs about quality are impacted by the size of the choice set. Allowing $\pi(\delta_{ijt}) = \pi(\delta_{ijt}, |S_{it}|)$ allows for the unconditional probability of choice to be altered by the size of the assortment. Letting $\mu_{it} = \mu_{it}(|S_{it}|)$ allows the size of the assortment to impact all alternatives at once, including incumbent and previously purchased ones.

4.1.2 Comparison to Full Information Discrete Choice Demand

Unlike random utility models that produce discrete choice probabilities based on a random utility component (often, ϵ), choice in this model is stochastic because of the consumer's uncertainty about choice payoffs. The previous equation bears a close resemblance to the choice probabilities in a multinomial logit model of demand, but they differ in two important ways. First, this model does not have unobserved taste shocks ϵ; its randomness in choice comes from consumers' incomplete information. This is beneficial to my setting where the size of the choice set is very large, because it partially addresses the problem highlighted in Ackerberg and Rysman (2005). In this model, new products do not introduce a new, equally valued attribute that differentiates the new product. Additionally, given the size of the choice set, it is unreasonable to assume consumers have full information about hundreds of products. Moreover, the reduced form findings—that purchase frequency declines in the size of the choice set —cannot be rationalized by a full information model of consumer choice.

Second, the form of this model is equivalent in prediction of behavior to a particular parameterization of a logit model of demand. Existing logit models (e.g., Draganska and Jain (2005); Ershov (2022)) include a congestion term that may account for this level shift, but they cannot account for the changes due to attention cost in sensitivity to consumption utility. My approach allows for a structural interpretation of this congestion term. As the cost of information rises, this model predicts consumers become less sensitive on average to price and quality information that requires attention or search. My approach provides a justification for the inclusion of the congestion term with a specific interpretation.

4.2 Implementation and Estimation

I estimate this model on a subset of consumers[24] from the consumer panel. To identify separately the parameters of interest, I parameterize the model as follows:

$$u_{ijt} = 1 + x_{ijt}\beta_i$$
$$\mu_{it} \propto exp(w'_{it}\theta_i)$$
$$\pi(\delta_{ijt}) \propto exp(d'_{ijt}\gamma_i)$$

Fixing the intercept in utility u allows for the multiplicative identification with the inverse of the information cost μ and an intercept term for inside options in π. I normalize the outsize option to have $u_0 = 0$ and $d_0 = 0$. I allow for intercept level shifts in utility through d and discrete x attributes (in this case, cuisine). I assume the cost of attention is weakly positive. This parameterization yields choice probabilities:

$$P_{ijt} = \frac{exp(d'_{ijt}\gamma_i + exp(w'_{it}\theta_i)(1 + x'_{ijt}\beta_i))}{1 + \sum_{k \in S_{it}} exp(d'_{ikt}\gamma_i + exp(w'_{it}\theta_i)(1 + x'_{ikt}\beta_i))} \tag{3}$$

	Included Covariates
w_{it}	Assortment Size, Unemployment Rate, Indicator for platform experience in the last 2 months
d_{ijt}	Assortment Size (separately for previously visited vs unvisited restaurants), # platform and restaurant visits, proxy for on-platform ads, # of Yelp reviews, Time FE
x_{ijt}	Price (post-tax, delivery inclusive), Cuisine, Delivery Distance

It is infeasible to identify a rich set of parameters as consumer-level fixed effects given the sparsity of choices and large number of restaurants. To flexibly accommodate heterogeneity in tastes and beliefs, I adapt the method of Bonhomme and Manresa (2015) and Bonhomme et al. (2021) for discrete consumer types using subsampling inference. In their method, continuous heterogeneity in unit-specific parameters is approximated in short panels by using observable variation in covariates, outcomes, and supplemental moments to cluster units into segments with shared parameters. My approach is similar: I segment consumers before estimation with k-means clustering[25] using

[24] I restrict the panel to consumers who are never mobile in their choices. I also consider only consumers who purchase at most once per week, since this model definition is premised on weekly discrete choice.

[25] Results are robust to alternative numbers of clusters.

consumer demographics, location, and average order frequency. Within a segment, parameters are estimated via maximum likelihood.

My inferential approach departs from Bonhomme et al. (2021) by simultaneously characterizing grouping uncertainty (into segments) and parameter uncertainty through subsampling the entire procedure (Romano et al., 2012).[26] Additionally, I do not have any shared parameters across segments/clusters. I sample a subset of users from the panel 1000 times before clustering, estimating, and conducting any counterfactuals.

Estimation and Counterfactual Simulation

---

for subsample `s in 1:1000` **do**
- Draw a random 98% subsample of consumers to form panel
- Cluster subsample into 10 groups. K-means clusters are based on supplementary demographic moments, location, and average consumption
- **for** cluster c_s `in 1:10` **do**
  - Estimate via MLE $(\hat{\gamma}_{is}, \hat{\theta}_{is}, \hat{\beta}_{is}) = (\hat{\gamma}_{cs}, \hat{\theta}_{cs}, \hat{\beta}_{cs}) \forall i \in c_s$
  - Compute elasticities and partial elasticities
- **end for**
- Simulate demand under alternative assortment strategies, varying maximum size and targeting schema (detailed in section 4.5)

end for

Summarize parameter values across iterations for each individual: **distribution of** $(\hat{\gamma}_{is}, \hat{\theta}_{is}, \hat{\beta}_{is})$

4.3 Identification

A key challenge in identifying this model is the separation of varying consumer prior beliefs, information costs, and taste heterogeneity. I assume that unobservable heterogeneity in tastes, information costs, and belief formation can be represented by a low dimensional factor representation, so that unobservable differences across consumers are well captured by the discrete, observable heterogeneity approach. Further discussion will explain identification *within* consumer or consumer type.

Identifying preferences for restaurant attributes (β), conditional on the rest of the model, is similar to most revealed preference discrete choice identification—it leverages cross-sectional and intertemporal changes to choice shares in response to differences and changes in x, holding the size of the choice set fixed. In the case of price sensitivity, I do not explicitly control for price endogeneity and rely on across-restaurant variation in prices and changes to municipal sales taxes over time.

[26] An additional benefit of this approach is to simplify significantly the representation of uncertainty in counterfactual exercises.

Prior beliefs $d\gamma$ and information costs $1/\mu$ are challenging to separate. When the assortment grows, it may affect expectations (which affects $\pi(\delta)$ and the information acquisition strategy) or the information costs (which affects $\pi(\delta)$ and the information acquisition strategy). Given the choice probability function derived from the assumptions about information costs, mechanically these two components affect the vector of product shares differently. If the assortment growth alters information costs, it changes not only the total probability of inside consumption but also its composition. For example, relative choice shares between two incumbent restaurants for which the consumer has equivalent information d will shift if unit information costs increase with the assortment size. One dimension on which we would expect to see such an effect is relative price sensitivity—if consumers are less price-sensitive when the assortment grows (though their underlying preference for price is unchanged by assumption), this would identify the effect of $|S_{it}|$ on μ. If, instead, the relative choice shares of these "ex ante identical" restaurants are unchanged, any changes to the total inside share would identify lower prior beliefs for inside incumbent firms.

There is additional parameterized variation in beliefs, information costs, and preferences outside of changes to the assortment. Differences in beliefs are identified from changes to choice probabilities across restaurants correlated with past order status changes while attributes x and information cost shifters w are fixed, for example. The effect of time-fixed belief shifters d, such as a proxy for on-platform prominence, are identified from cross sectional variation in choice shares conditional on x and w. Differences in baseline information costs $1/\mu$ within a consumer type are identified from differences in relative sensitivity to x, holding fixed d—by assumption, consumption utility is the same function of x for all consumers of a type. Much of this variation is cross-sectional, similar to interacting consumer demographics with preference parameters. Finally, I identify preference parameters β for attributes that require attention (defined by assumption) and the remainder of γ from cross-sectional and temporal changes to choice shares in response to x and d, holding the size of the choice set and determinants of attention costs fixed. In the case of identifying price sensitivity, I rely mostly on across-restaurant variation in prices, though changes to municipal sales taxes allow me to use some variation over time in prices.

4.4 Estimation Results

I estimate three specifications that vary in terms of how the size of the choice set impacts consumers. I test a version of the model where the size of the choice set impacts expectations based on δ alone, where it impacts information costs μ alone, and where it impacts both. The size of the choice set does not significantly impact the information cost in either specification where this is possible. The size of the choice set does significantly impact prior beliefs about utility for restaurants that the consumer has not tried in the past, though not for restaurants that they have consumed. Specification tests reject the version of the model where the size of the choice set enters the information cost alone, and the version of the model where $|S|$ effects both μ and δ. I explore counterfactual assortment strategies using the model where only consumer prior beliefs about products are a function of the choice set size.

Table 11: Elasticity Estimates

	Assortment Size on No-Purch	Own-Price	Restaurant Distance
Mean Household	0.02128	-2.703	-0.290
Median Household	0.00785	-2.531	-0.267
Variance of Means	0.00394	4.507	0.101
Mean Household Variance	0.34979	146.521	4.293
Median Household Variance	0.00002	0.790	0.048

The main model effect of interest is how choice probabilities change with assortment expansion. This partial elasticity is not equivalent to a full counterfactual of removing choices. Instead, it tells us how the probability of the no purchase changes as expectations and/or information costs adjust to the assortment size, holding the real size of the set and its contents fixed.[27] For most consumers, this elasticity is weakly positive. However, for many consumers this elasticity is very close to zero—there is heterogeneity in how much users are negatively impacted by larger assortments. Figure 1 shows the distribution of this semi-elasticity across consumers. The average consumer has an partial-elasticity of 0.02, and the median consumer's partial-elasticity is 0.01. Consumers are sensitive to the assortment size, but the effect is very small. I report further summary statistics for assortment size, price, and restaurant distance elasticities in Table 11. Full parameter estimates are available in Appendix A2.

While the estimates suggest that consumers form expectations about their net returns to attention and consumption based on the size of this choice set, I don't explicitly model this connection. I include a linear term for the size of the choice set (interacted with restaurant experience), but future work should explore how to explicitly empirically model consumers' expectations of each product in a large assortment.

4.5 Counterfactual Assortment Reductions

Given the reduced form and structural evidence for a penalty to large choice sets, I conduct counterfactual restrictions of each individual's choice set. Unlike in most demand models (e.g.

[27] Specifically, I am interested in the distribution $\frac{\partial P_{i0t}}{\partial |S_{it}|} \frac{|S_{it}|}{P_{i0t}}$

Figure 1: Assortment Size Semi-Elasticity of No-Purchase

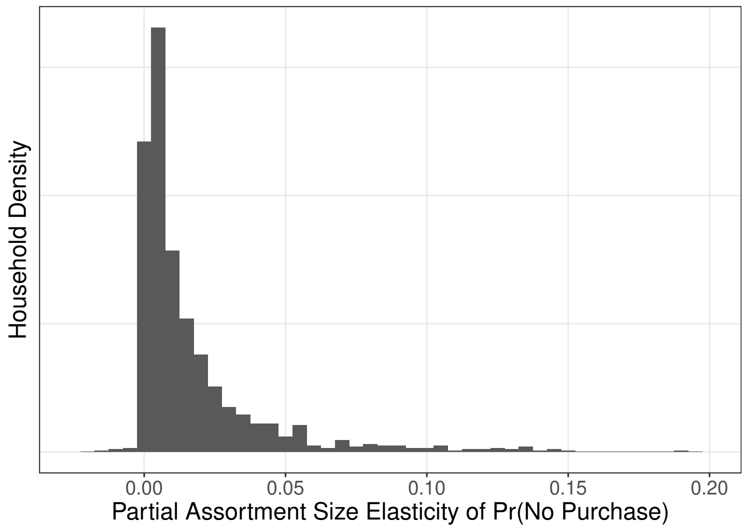

removing options in a logit demand system holding attributes fixed), the direction of these counterfactual results are not ex-ante clear and therefore worth testing empirically. Whether reducing the size of the assortment benefits the probability of purchase depends two factors: how beneficial an option is in terms of consumption, and how large the penalty is relative to this benefit.

I use several targeting metrics to test assortment size reductions for each individual. I target which restaurants to remove based on the consumer choice model. I perform this targeting based on three model-driven metrics: expected revenue ($P_{ijt} * r_{ijt}$), choice probability (P_{ijt}), and consumption utility $u_{ijt} = 1 + x_{ijt}\beta_i$ (which does not incorporate contextual factors like choice history and assortment size). The expected revenue metric reflects the platform's objective function, where a commission is earned on each sale. The expected revenue from an order varies based on the average basket size at the restaurant. I compare these targeting metrics to reducing the assortment size by removing alternatives at random.

For each targeting metric, I simulate a series of assortment reductions for each individual. The counterfactual is evaluated at each assortment size as if all consumers cannot have more than the proposed number of restaurants, i.e. the assortment size is a cap. Consumers may still face smaller choice sets if their area has fewer restaurants supplied than the targeted size, but no restaurants are added.[28]

Removing choices alters the distribution of match values/consumption utilities and the con-

[28]To keep these evaluations tied to the variation in the data, I restrict the maximum assortment size caps to start above 30 restaurants. I test variations smaller than 30 if a consumer has fewer than 30 restaurants in the data.

sumer's prior choice probabilities. The first mechanic is captured by many demand models. If the removed restaurants were unlikely to be chosen, this may minimally impact their probability of any inside purchase. The second mechanic is what allows the model to improve the inside share with a smaller choice set. If the pre-attention/unconditional choice probabilities decrease with the size of the choice set, consumers will be less likely to purchase any product.

Figure 2 plots the average normalized change in revenue by targeting metric for a sample week if all consumers face the same maximum assortment size across these caps. Reducing assortments improves the platform's revenue from existing consumers when the reductions are targeted using consumers' choice histories (quantity or revenue prediction targeting), but not when they are randomized or attribute targeted, shown in Table 12. This contrast highlights the interaction between the size of the assortment and its contents. Reducing the assortment can increase weekly revenues 13% if the reductions favor likely-to-be-chosen restaurants. However, this restriction is imposed uniformly on customers, which hides that the best-case assortment size could vary considerably across customers. I contrast these platform-uniform restrictions with offering each individual

Figure 2: Counterfactual Revenues: Uniform Assortment Size Maximum

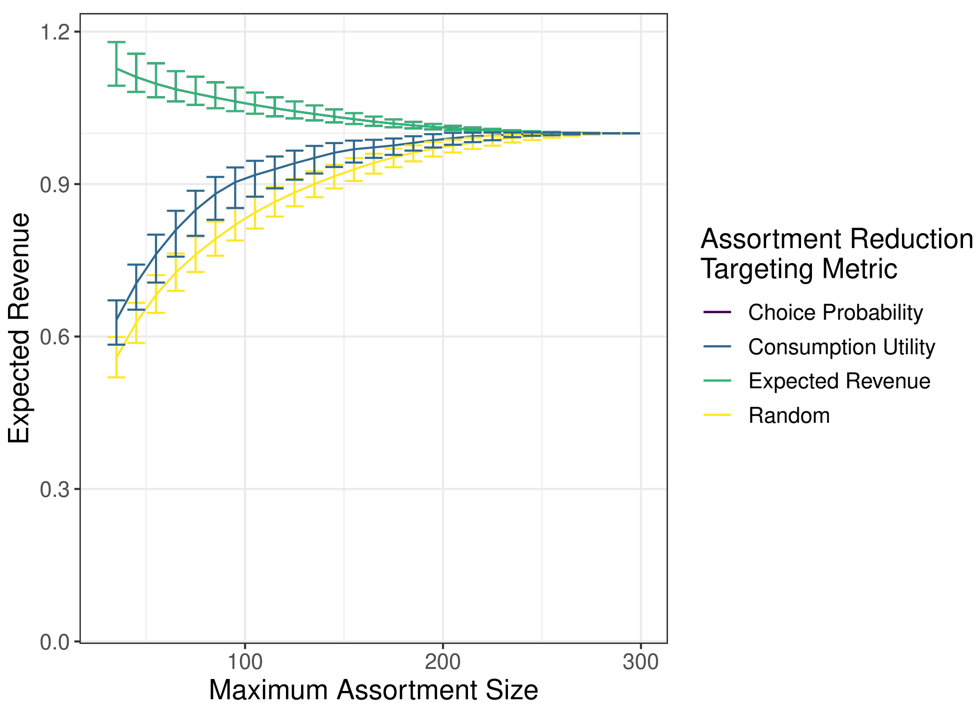

Note: This takes the average normalized revenue for the assortment size cap across subsampling iterations; error bars show 2.5 and 97.5 quantiles. Choice Probability and Expected Revenue lines are indistinguishable.

consumer in the data their own targeted maximum assortment size and contents. Targeted assort-

Table 12: Uniform Assortment Size Revenues

Targeting Metric	Mean Improvement	Median Improvement	2.5 Pctile	97.5 Pctile
Choice Probability	1.1274	1.1231	1.0937	1.1796
Random	1.0000	1.0000	1.0000	1.0000
Expected Revenue	1.1275	1.1232	1.0937	1.1794
Consumption Utility	1.0001	1.0003	0.9992	1.0005
		Results reported as a ratio with base of current assortment		

Table 13: Uniform Assortment Size Sales

Targeting Metric	Mean Improvement	Median Improvement	2.5 Pctile	97.5 Pctile
Choice Probability	1.1248	1.1202	1.0917	1.1776
Random	1.0000	1.0000	1.0000	1.0000
Expected Revenue	1.1248	1.1202	1.0918	1.1776
Consumption Utility	1.0001	1.0002	0.9993	1.0004
		Results reported as a ratio with base of current assortment		

ment sizes *and* contents improve weekly revenue by 22%. Many targeted choice set sizes are the same as a uniform cap. The distribution of these assortment size maxima are plotted by targeting schema in Figure 4. Under this model, a minimal assortment size maximizes revenue or sales quantity for most consumers. However, if the firm is restricted to using only restaurant attributes to target choice set contents, using the individually targeted choice set size improves expected revenue significantly (versus in the uniform case, where it does not).

Tables 13 and 15 show the expected increase in weekly order quantities, which in the both uniform and targeted assortment sizes increase 13 to 17% in the best case. I contrast this to revenue improvement (13 to 22%) to point out that most of the gain is getting consumers to order more often. However, the platform has an incentive to target the assortments towards higher-platform-margin restaurants when targeting is done by expected revenue. Targeting by choice probability (expected quantity) produces similar quantity improvements, but it does not capture as much platform revenue.

Figure 5 shows the distribution of the optimal assortment restriction to the individual's baseline assortment. Since my approach respects the supply of restaurants on the platform, I cannot rule out that some consumers would purchase more if supplied with larger assortments. For most consumers, this represents a considerable reduction in the realized choice—removing over 50% of the existing restaurants.

The optimal choice set sizes discussed above vary across consumer segments in terms of the flatness of the firms' objective function. For most consumers, removing alternatives improves

Figure 3: Counterfactual Sales: Uniform Assortment Size Maximum

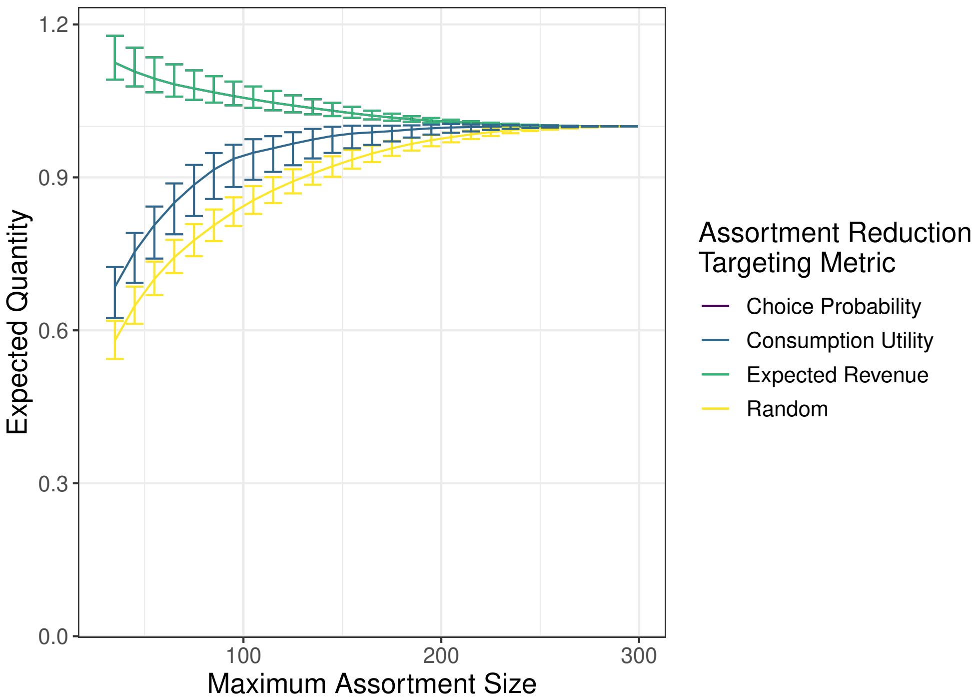

Note: This takes the average normalized sales for the assortment size cap across subsampling iterations; error bars show 2.5 and 97.5 quantiles. Choice Probability and Expected Revenue lines are indistinguishable.

Figure 4: Distribution of Individual-Specific Assortment Size Caps

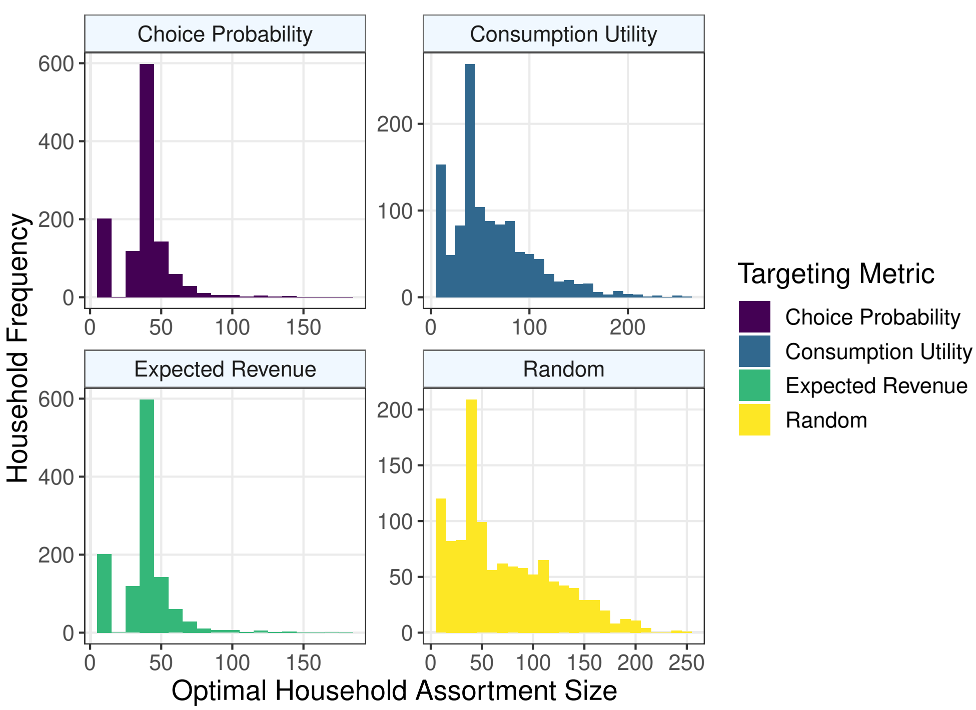

Table 14: Targeted Assortment Size Revenues

Targeting Metric	Mean Improvement	Median Improvement	2.5 Pctile	97.5 Pctile
Choice Probability	1.2155	1.1579	1.1130	1.4214
Random	1.0693	1.0527	1.0377	1.1818
Expected Revenue	1.2161	1.1583	1.1139	1.4218
Consumption Utility	1.1580	1.0873	1.0565	1.3853

Results reported as a ratio with base of current assortment

Table 15: Targeted Assortment Size Sales

Targeting Metric	Mean Improvement	Median Improvement	2.5 Pctile	97.5 Pctile
Choice Probability	1.1694	1.1555	1.1103	1.2603
Random	1.0522	1.0480	1.0322	1.0945
Expected Revenue	1.1690	1.1550	1.1099	1.2597
Consumption Utility	1.1070	1.0865	1.0548	1.1981

Results reported as a ratio with base of current assortment

purchase probability and revenue, but the shape is not as pronounced as in the aggregate total (Figure 4). In other words, the firm's objective function is fairly flat for some consumers, but not for others. Future work should explore how well this out-of-sample prediction (large reductions in each individual's choice set).

Figure 5: Distribution of Individual-Specific Assortment Size Reductions

The structure of this counterfactual is a partial-equilibrium concept; it does not consider how restaurants might react to restricting individual choice sets. However, since the platform is restricting the size of the potential consumer base by pruning low-purchase-likelihood consumers, this should not be too detrimental to restaurants. Moreover, since the restriction improves the probability of purchase, this could net-benefit many restaurants. I do not explore the compositional effect on restaurants further, because this estimation exercise is focused only on a small subset of consumers in one area. Future tests should consider how restaurants might react in equilibrium (for example, whether this generates market power for restaurants).

The counterfactuals are stylized versions of potential improvements that can be made by platforms. Platforms may instead consider heavy personalization as an alternative (Donnelley et al., 2021). The platform does not need to prevent their existing customers from ever accessing all the restaurants that serve their location. In practice, platforms may choose to allow consumers to find any restaurant in their set if they search for it directly.

5 Conclusion and Future Work

In this paper, I document that 'more is not better' for some consumers - particularly, those who already have participated in a category. Unlike previous work, I show this effect in large, real-world assortments over the consumer lifetime. In an online food delivery platform, larger product assortments—more restaurants—drive increased consumer adoption of the platform, but lower the rate at which existing consumers order. This effect cannot be rationalized by most demand models. I use a model of attention allocation where the size of the choice set impacts consumers' beliefs and attention costs to show how intermediaries can reduce the assortment for each individual in a targeted manner.

The online restaurant delivery market provides an ideal lab for isolating the effect of assortment size, since assortments observably vary frequently and across individuals. However, by construction, choice in this market is always discrete—consumers only order from one restaurant at a time. The discreteness of the choice allows me to identify choice frictions from larger assortments among repeat customers, but the net negative effect may not generalize to basket situations. For example, in grocery retailing or other markets where the norm is baskets containing multiple categories and multiple products within category, the benefit of variety may outweigh its costs. This could explain differences between the findings in this paper and in related work in grocery retailing (Borle et al., 2005). Still, there are many markets where discrete choice is relevant, and choice frictions may dominate. In infrequent, large-ticket categories (computers, cars), these choice frictions from variety may be hard to measure, but findings in this work can shed light on the potential drawbacks to more variety in these markets.

Future work on this topic falls into three groups. First, the findings from this paper can be tested by platforms. In my counterfactual exercises, I find that using imperfectly targeted measures to reduce the choice set can still produce gains. Platforms' internal recommendations model can be tested as ways to reduce assortments. Second, the results here can be extended to include menu-item level analysis if such data were available. Third, more detailed consumer search data can be leveraged to better understand the exact mechanism for these choice frictions in large assortments. Exploring this issue further will help determine where, when, and how assortment reductions should be implemented.

References

Abadie, Alberto, Susan Athey, Guido Imbens, and Jeffrey Wooldridge, "When Should You Adjust Standard Errors for Clustering?," 2017.

Ackerberg, Daniel and Marc Rysman, "Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects," *RAND Journal of Economics*, 2005, pp. 771–788.

Aguiar, Luis, Joel Waldfogel, and Sarah Waldfogel, "Playlisting Favorites: Measuring Platform Bias in the Music Industry," *International Journal of Industrial Organization*, 2021, p. 102765.

Almagro, Milena and Tomás Domınguez-Iino, "Location sorting and endogenous amenities: Evidence from amsterdam," 2021. Working Paper.

Berry, Steven, Alon Eizenberg, and Joel Waldfogel, "Optimal Product Variety in Radio Markets," *RAND Journal of Economics*, 2016, *47* (3), 463–497.

Berry, Steven T. and Joel Waldfogel, "Free Entry and Social Inefficiency in Radio Broadcasting," *The RAND Journal of Economics*, 1999, *30* (3), 397.

Boatwright, Peter and Joseph C. Nunes, "Reducing assortment: An attribute-based approach," *Journal of Marketing*, 2001.

Bonhomme, Stéphane and Elena Manresa, "Grouped Patterns of Heterogeneity in Panel Data," *Econometrica*, 2015, *83* (3), 1147–1184.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa, "Discretizing Unobserved Heterogeneity," *Econometrica*, 2021. Forthcoming.

Borle, Sharad, Peter Boatwright, Joseph B. Kadane, Joseph C. Nunes, and Galit Shmueli, "The effect of product assortment changes on customer retention," *Marketing Science*, 2005.

Briesch, Richard A., Pradeep K. Chintagunta, and Edward J. Fox, "How does assortment affect grocery store choice?," *Journal of Marketing Research*, 2009.

Broniarczyk, Susan M., Wayne D. Hoyer, and Leigh McAlister, "Consumers' Perceptions of the Assortment Offered in a Grocery Category: The Impact of Item Reduction," *Journal of Marketing Research*, 1998.

Bronnenberg, Bart J., "The provision of convenience and variety by the market," *RAND Journal of Economics*, 2015.

Callaway, Brantly and Pedro HC Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2020.

Caplin, Andrew and Mark Dean, "Revealed preference, rational inattention, and costly information acquisition," *American Economic Review*, 2015.

Casadesus-Masanell, Ramon and Hanna Halaburda, "When does a platform create value by limiting choice?," *Journal of Economics & Management Strategy*, 2014, *23* (2), 259–293.

Chade, Hector and Lones Smith, "Simultaneous search," *Econometrica*, 2006, *74* (5), 1293–1307.

Chaisemartin, Clément De and Xavier d'Haultfoeuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–96.

Chandon, Pierre, J. Wesley Hutchinson, Eric T Bradlow, and Scott H Young, "Does In-Store Marketing Work? Effects of the Number and Position of Shelf Facings on Brand Attention and Evaluation at the Point of Purchase," *Journal of Marketing*, 2009, *73* (6), 1–17.

Chernev, Alexander, Ulf Böckenholt, and Joseph Goodman, "Choice overload: A conceptual review and meta-analysis," 2012.

Chu, Junhong and Puneet Manchanda, "Quantifying Cross and Direct Network Effects in Online Consumer-to-Consumer Platforms," *Marketing Science*, 2016, *35* (6), 870–893.

Dixit, Avinash K. and Joseph E. Stiglitz, "Monopolistic competition and optimum product diversity," *The American Economic Review*, 1977, *67*, 297–308.

Donnelley, Robert, Ayush Kanodia, and Ilya Morozov, "The Long Tail Effect of Personalized Rankings," *Working Paper*, 2021.

Draganska, Michaela and Dipak C. Jain, "Product-line length as a competitive tool," *Journal of Economics and Management Strategy*, 2005.

Drèze, Xavier, Stephen J. Hoch, and Mary E. Purk, "Shelf management and space elasticity," *Journal of Retailing*, 1994.

Eisend, Martin, "Shelf space elasticity: A meta-analysis," *Journal of Retailing*, 2014, *90* (2), 168–181.

Ershov, Daniel, "Competing with Superstars in the Mobile App Market," 2022. Working Paper.

Farronato, Chiara, Jessica Fong, and Andrey Fradkin, "Dog Eat Dog: Measuring Network Effects Using a Digital Platform Merger," *Management Science*, 2022. Forthcoming.

Gu, Chris and Yike Wang, "Consumer online search with partially revealed information," *Management Science*, 2022, *68* (6), 4215–4235.

Halaburda, Hanna, Mikołaj Jan Piskorski, and Pinar N. Yildirim, "Competing by restricting choice: The case of matching platforms," *Management Science*, 2018.

Hotelling, Harold, "Stability in Competition," *The Economic Journal*, 1929.

Illanes, Gastón and Sarah Moshary, "Market Structure and Product Assortment: Evidence from a Natural Experiment in Liquor Licensure," Working Paper 27016, National Bureau of Economic Research April 2020.

Iyengar, Sheena S. and Mark R. Lepper, "When choice is demotivating: Can one desire too much of a good thing?," *Journal of Personality and Social Psychology*, 2000.

Joo, Joonhwi, "Rational Inattention as an Empirical Framework: Application to the Welfare Effects of New Product Introduction," *Journal of Marketing Research*, 2022.

Kahn, Barbara E. and Brian Wansink, "The Influence of Assortment Structure on Perceived Variety and Consumption Quantities," *Journal of Consumer Research*, 2004.

Kamenica, Emir, "Contextual inference in markets: On the informational content of product lines," *American Economic Review*, 2008.

Kim, Yewon, "Consumer Retention under Imperfect Information," 2021. Working Paper.

Kuksov, Dmitri and J. Miguel Villas-Boas, "When More Alternatives Lead to Less Choice," *Marketing Science*, 2010.

Lancaster, Kelvin, "Socially optimal product differentiation," *American Economic Review*, 1975.

_ , "The Economics of Product Variety: A Survey," *Marketing Science*, 1990.

Li, Jun and Serguei Netessine, "Higher market thickness reduces matching rate in online platforms: Evidence from a quasiexperiment," *Management Science*, 2020, *66* (1), 271–289.

Lin, Xiliang, "Disaggregate network effects on two-sided platforms." PhD dissertation, The University of Chicago 2017.

Matêjka, Filip and Alisdair McKay, "Rational inattention to discrete choices: A new foundation for the multinomial logit model," *American Economic Review*, 2015.

Nocke, Volker and Patrick Rey, "Consumer Search and Choice Overload," 2022. Working Paper.

Quan, Thomas W. and Kevin R. Williams, "Product variety, across-market demand heterogeneity, and the value of online retail," *RAND Journal of Economics*, 2018.

Reshef, Oren, "Smaller Slices of a Growing Pie: The Effects of Entry in Platform Markets," 2022. Working Paper.

Romano, Joseph P, Azeem M Shaikh et al., "On the uniform asymptotic validity of subsampling and the bootstrap," *The Annals of Statistics*, 2012, *40* (6), 2798–2822.

Scheibehenne, Benjamin, Rainer Greifeneder, and Peter M. Todd, "Can There Ever Be Too Many Options? A Meta-Analytic Review of Choice Overload," *Journal of Consumer Research*, 2010, *37* (3), 409–425.

Sims, Christopher A, "Implications of rational inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

Sloot, Laurens M, Dennis Fok, and Peter C Verhoef, "The short-and long-term impact of an assortment reduction on category sales," *Journal of Marketing Research*, 2006, *43* (4), 536–548.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2020.

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

A1 Data Summary

Table A1: Summary Statistics: Consumer Panel

Statistic	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)
Orders	19.85	31.77	2	3	8	22
Restaurants Tried	7.15	8.63	1	2	4	9
Spending	519.54	885.53	14.70	77.90	193.84	553.68
Cuisines Tried	4.15	3.16	1	2	3	6

Figure A6: Census Tract Panel Choice Set Sizes

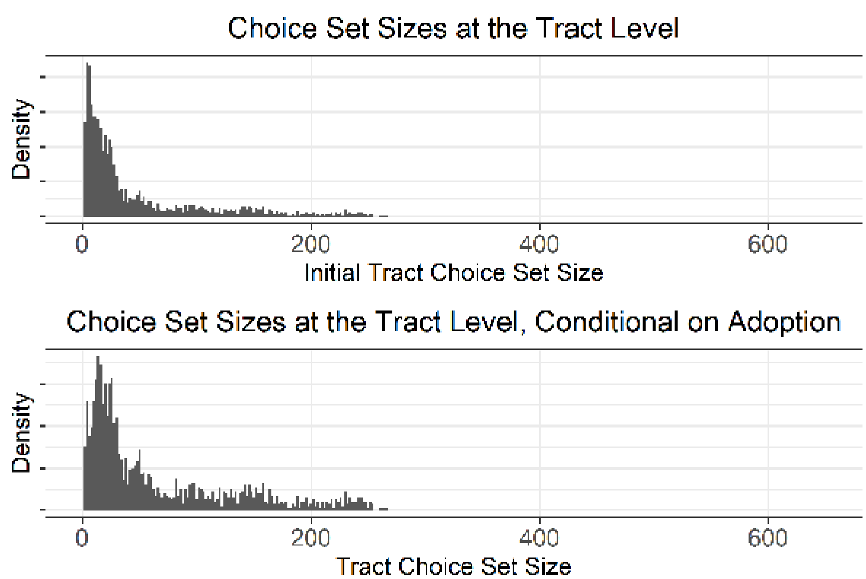

Figure A7: Individual Panel Choice Set Sizes at Start

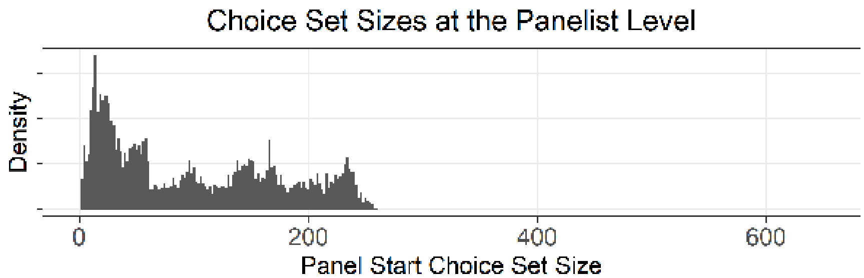

Figure A8: Individual Panel Choice Set Sizes at End

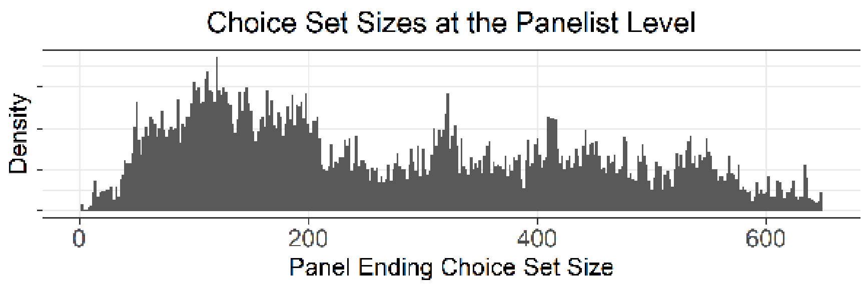

A2 Selected Estimation Results

I reported the distribution of elasticity estimates for 3 measures of interest in Table 11: the effect of assortment size on no-purchase, own-price elasticity, and the distance elasticity of restaurant-consumer matches. I also report summary statistics for all parameters (excluding time effects) in Table A2.

I plot the distribution across households of the mean parameter estimates for each household for selected elements of γ: the coefficients on assortment size of tried and untried restaurant counts in Figure A9. Note that the scale of these parameters differ, as do the covariates in the data. For most consumers, the number of untried restaurants is nearly the entire assortment, while the median consumer has tried 4 restaurants in the entire sample.

I also plot the responsiveness of consumers to prices (Figure A10) and restaurant distance (Figure A11) in elasticity terms. The median consumer has an own-price elasticity of -2.5, and 2.8% of consumers have inelastic demand. Most households have a disutility of distance to the restaurant.

Figure A9: Effect of Assortment Size on Consumer's Subjective Beliefs

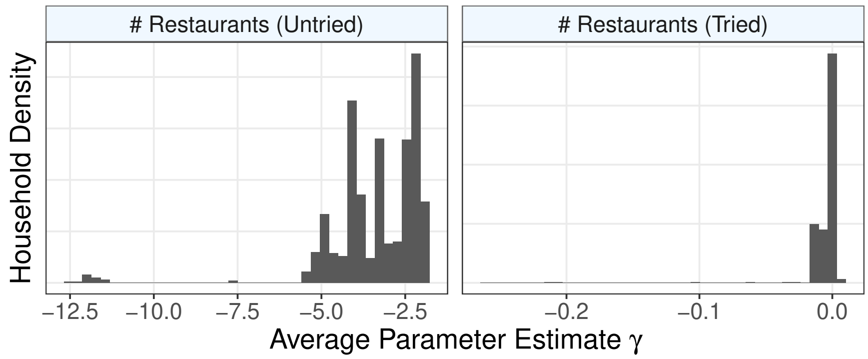

Figure A10: Own-Price Elasticity across Households

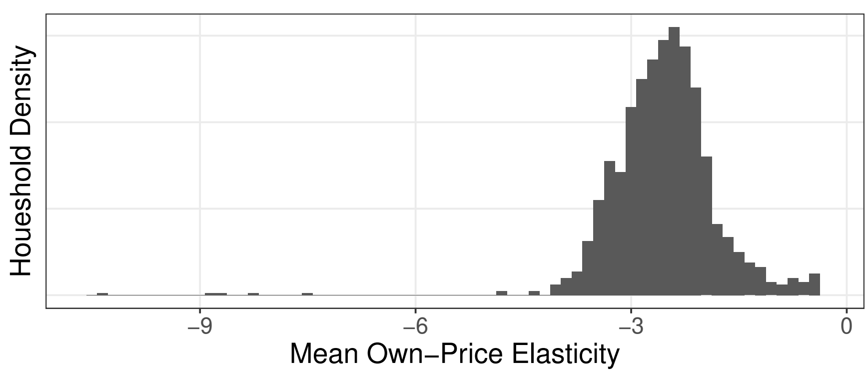

Figure A11: Restaurant Distance Elasticity across Households

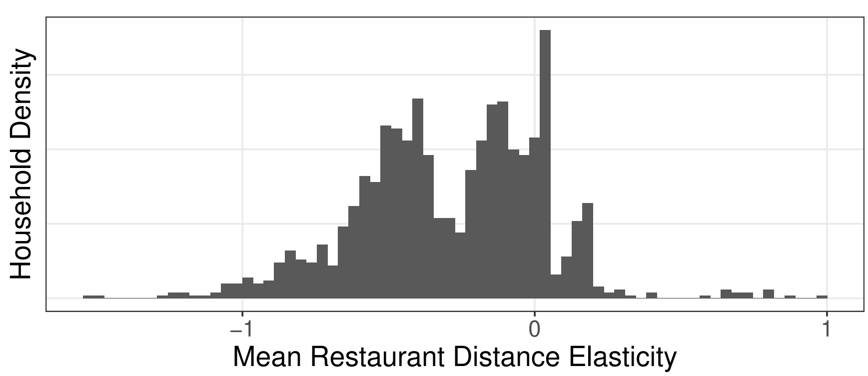

Table A2: Parameter Estimates

Variable	Mean	Median	Var	Q(2.5)	Q(97.5)	Parameter
American	-45.841	-1.039	39339787.0	-3.673	0.117	β
Asian	5.329	-0.408	531214.9	-14.241	1.422	β
Breakfast	14.799	-0.304	4247365.6	-2.795	1.983	β
Chinese	-15.453	-0.591	3544416.2	-3.539	0.646	β
Indian	0.430	0.215	1808034.8	-3.757	3.565	β
Italian	8.325	-0.902	36189740.5	-3.318	0.080	β
Japanese	-7.606	-0.521	479817.9	-4.384	3.115	β
Juice Bars Smoothies	0.874	0.497	22217.6	-33.990	48.497	β
Mediterranean	-7.065	-0.364	4363193.0	-9.683	3.227	β
Mexican	-3.940	-0.884	264884.4	-19.573	4.113	β
Noodles	-2.952	-0.671	379430.9	-3.855	0.987	β
Pizza	-5.287	-1.229	17195.9	-8.210	-0.037	β
Price (USD)	2.378	-0.091	124840.8	-0.176	-0.002	β
Rest Distance (km)	-18.611	-0.086	12485867.8	-0.523	0.236	β
Salads	2.795	0.012	20311.7	-11.794	19.147	β
Sandwiches	-41.684	-8.733	8553.9	-365.599	112.928	β
Prev Rest Orders	0.078	0.075	0.4	0.036	0.112	γ
Prev Orders	0.263	0.133	40.9	0.025	0.371	γ
Platform Ad Proxy	-0.053	-0.038	0.6	-0.127	-0.004	γ
Yelp Review Count	-0.000	-0.000	0.0	-0.000	-0.000	γ
# Untried Restaurants	-9.757	-1.623	1942107.7	-15.619	-0.349	γ
# Tried Restaurants	-0.006	-0.002	0.0	-0.015	0.000	γ
Ordered Within 2 Mos	-29.562	-2.895	4832285.2	-55.056	1.518	μ
Unemployment Rate	-0.412	-0.013	8.2	-3.805	0.030	μ

ONLINE APPENDICES

OA1 Additional Tables

Table OA1 separates treatment by whether the restaurant belongs to a chain. Chain restaurants are defined to include large, national quick-serve and fast casual restaurants, regional chains, and local chains with at least 5 outlets. The estimated effect of the entry of an independent restaurant is very similar to the main effect - the weekly orders and spending decline when an independent restaurant enters. However, results are mixed for chain entry - orders and spending may actually rise with the entry of chain restaurants onto the platform. Note that chains comprise 11% of restaurants on the platform.

Table OA2 reports the result of a regression on time between orders for the sample of consumers with search data. These regressions condition on weeks with positive consumption, and regresses the weeks until the next order on the size of the choice set at the time of the present order,

Table OA1: Effect of Chain Restaurants on Returning Customers

	Dependent variable:	
	Weekly Orders	Weekly Spending (USD)
	(1)	(2)
Independent Restaurant Count	-0.0001^{**}	-0.001
	(0.00005)	(0.001)
Chain Restaurant Count	0.0003^{*}	0.009
	(0.0001)	(0.004)
Observations	2,016,651	2,016,651
R^2	0.240	0.223
Adjusted R^2	0.210	0.192
Note:		*p<0.05; **p<0.01; ***p<0.001 All specs include ZCTA-Week and Individual FEs Standard Errors clustered at ZCTA-week level

Figure OA12: Average Interpurchase Time by Search Duration

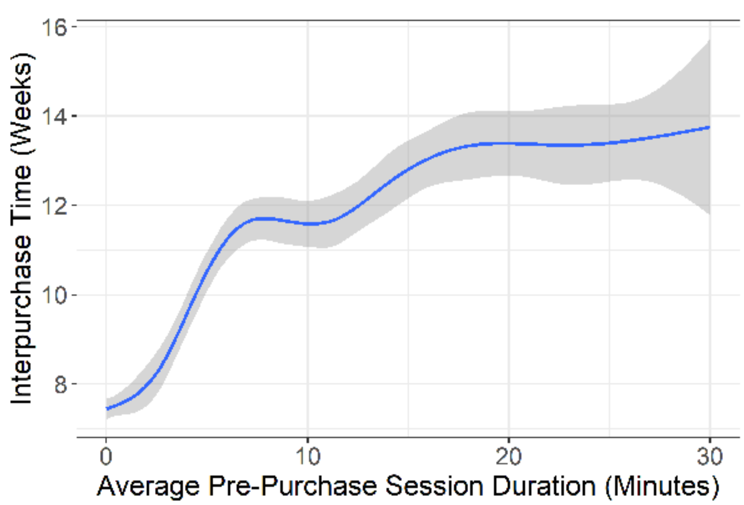

controlling for the same fixed effects as in the main regressions. This does not control for selection into purchase, but, conditional on having made a purchase, larger assortments are correlated with a longer wait until the subsequent purchase. This could be a form of learning about the pain of search/choice through experience. New customers, who join at higher rates due to assortment growth, may not have yet experienced such costs.

Table OA2: Assortment Size and Time until next Purchase

	Dependent variable:	
	Interpurchase Time (Weeks)	
	(1)	(2)
Restaurant Count	0.013	0.013
	(0.007)	(0.007)
Search Duration (Mins)		0.013*
		(0.005)
Observations	100,310	100,310
R^2	0.580	0.582
Adjusted R^2	0.323	0.321
Note:		*p<0.05; **p<0.01; ***p<0.001
		All specs include Individual and ZCTA-Week FEs
		Standard Errors clustered at ZCTA-week level

Assortment growth may have heterogeneous effects across consumers. Consumers may value variety differently, and they may realize any costs of sifting through many products differently. I allow the effect of assortment size to differ by the degree of observed variety consumption in the panel.[29] I construct a ratio of each consumer's total variety consumed to total consumption, and interact this ratio with the size of the assortment. Table OA4 shows these estimates. Consumers who maximally vary their consumption - e.g. ordering from a different restaurant every time - order slightly more when variety on the platform increases. In contrast, consumers who vary their consumption little, are more impacted negatively by assortment expansion. These effects persist even when excluding high frequency users, who make up a greater share of the low-ratio users.

An alternative measure of consumption variety would compare users in the top and bottom quartiles of variety consumption and run separate analysis for each. This partially conditions on order frequency - low variety users have a lower average number of orders. However, the results in this comparison are not consistent with the variety ratio results - here, it is users with a more

[29]These results condition on outcomes and should be taken as descriptive only.

Table OA3: Assortment Size and Search Duration by Purchase Type

	Dependent variable:		
	Search Duration (Minutes)		
	(1)	(2)	(3)
Rest Ct×New Purchase	0.059***	0.069**	0.062**
	(0.001)	(0.022)	(0.022)
Rest Ct×No Purchase	−0.001	−0.014	−0.015
	(0.001)	(0.022)	(0.022)
Rest Ct×Repeat Purchase	0.036***	0.032	0.031
	(0.001)	(0.022)	(0.022)
Conditions on Search	N	Y	Y
Selection Controls?	NA	N	Y
Observations	1,436,956	84,630	84,630
R^2	0.252	0.564	0.573
Adjusted R^2	0.212	0.239	0.256

Note: *p<0.05; **p<0.01; ***p<0.001
All specs include ZCTA-Week and Individual FEs
Omits search selection first stage residual control
Standard Errors clustered at ZCTA-week level

Table OA4: Effect of Assortment Size on Weekly Orders by User Consumption Variety

	Dependent variable:	
	Weekly Orders	
	(1)	(2)
Restaurant Count	−0.001***	−0.0005***
	(0.00004)	(0.00002)
Restaurant Count x Variety Ratio	0.001***	0.001***
	(0.00002)	(0.00001)
Exclude Highest 20% Frequency Users?	N	Y
Observations	2,058,406	1,640,421
R^2	0.240	0.109
Adjusted R^2	0.209	0.066

Note: *p<0.05; **p<0.01; ***p<0.001
All specs include Individual and ZCTA-Week FEs
Standard Errors clustered at ZCTA-Week level
Variety ratio of unique restaurants to total orders

varied basket who are more negatively impacted by assortment growth.

Table OA5: Effect of Assortment Size on Weekly Orders by User Consumption Variety

	Dependent variable:
	Weekly Orders
Panel A: High Variety Users	
Restaurant Count	−0.0001
	(0.0001)
Panel B: Low Variety Users	
Restaurant Count	0.0001***
	(0.00002)
Observations	529,361
R^2	0.256
Adjusted R^2	0.165
Note:	*p<0.05; **p<0.01; ***p<0.001
	All specs include Individual and ZCTA-Week FEs
	Standard Errors clustered at ZCTA-week level

Table OA6: Lagged and Contemporaneous Effects of Assortment Size on Spending

	Dependent variable:	
	Weekly Spend (USD)	
	(1)	(2)
Restaurant Count	0.001	
	(0.001)	
Restaurant Entry		−0.118***
		(0.012)
Lag Restaurant Count		0.002
		(0.001)
Observations	2,058,406	2,058,406
R^2	0.221	0.221
Adjusted R^2	0.190	0.190
Note:		*p<0.05; **p<0.01; ***p<0.001
		All specs include Individual and ZCTA-Week FEs
		Standard Errors clustered at ZCTA-Week level

Table OA7: Varying Marginal Effect of Assortment Size on Returning Customers

	Dependent variable:	
	Weekly Orders	Weekly Spending (USD)
	(1)	(2)
Restaurant Count × 1 to 140 Restaurants	−0.0001	0.001
	(0.00004)	(0.001)
Restaurant Count × 140 to 278 Restaurants	−0.0001	0.002
	(0.00003)	(0.001)
Restaurant Count × 278 to 417 Restaurants	−0.0001*	0.001
	(0.00003)	(0.001)
Restaurant Count × 417 to 555 Restaurants	−0.0001	0.002
	(0.00003)	(0.001)
Restaurant Count × 555 to 700 Restaurants	−0.0001*	0.0001
	(0.00003)	(0.001)
Observations	2,058,406	2,058,406
R^2	0.237	0.221
Adjusted R^2	0.207	0.190

Note: *p<0.05; **p<0.01; ***p<0.001
All specs include Individual and ZCTA-Week FEs
Standard Errors clustered at ZCTA-week level

Table OA8: Impact of Relevant Restaurant Entry on Returning Customers

	Dependent variable:	
	Weekly Spending	
	(1)	(2)
Ever Consumed Cuisine Restaurants	−0.020***	−0.012***
	(0.001)	(0.001)
Never Consumed Cuisine Restaurants		0.008***
		(0.001)
Observations	2,058,406	2,058,406
R^2	0.221	0.221
Adjusted R^2	0.190	0.190

Note: *p<0.05; **p<0.01; ***p<0.001
All specs include Individual and ZCTA-Week FEs
Standard Errors clustered at ZCTA-week level

To account for selection into search, I consider controlling for the residuals from a first stage linear probability model for whether the consumer engages in any search during the week. Controlling for selection in this manner does not alter the qualitative results for conditional search behavior.

Table OA9: Effect of Assortment Size on Search Duration and Purchase given Any Search

	Dependent variable:			
	Weekly Search Duration Conditional on Search		Weekly Orders Conditional on Search	
	(1)	(2)	(3)	(4)
Restaurant Count	0.020	0.008	−0.0005	−0.001
	(0.014)	(0.013)	(0.0004)	(0.0004)
Selection Controls?	N	Y	N	Y
Observations	135,414	135,414	135,414	135,414
R^2	0.488	0.530	0.521	0.522
Adjusted R^2	0.241	0.303	0.289	0.290

Note: *p<0.05; **p<0.01; ***p<0.001
All specs include ZCTA-Week and Individual FEs
Omits search selection controls
Standard Errors clustered at ZCTA-week level

Table OA10: Effect of Assortment Size on Weekly Search Sessions

	Dependent variable:		
	Weekly Sessions		
	(1)	(2)	(3)
Restaurant Count	−0.0002**	−0.0003***	−0.0002**
	(0.0001)	(0.0001)	(0.0001)
Cuisine Count		0.004***	
		(0.001)	
Cuisine Entropy			0.022
			(0.013)
Count Elasticity	-0.2403	-0.3212	-0.2364
Observations	1,436,956	1,436,956	1,436,956
R^2	0.262	0.262	0.262
Adjusted R^2	0.222	0.222	0.222

Note: *p<0.05; **p<0.01; ***p<0.001
All specs include ZCTA-Week and Individual FEs
Standard Errors clustered at ZCTA-week level

Table OA11: Local Effect of Assortment Size Conditional on Weekly Search

	Dependent variable:					
	Weekly Orders Given Search Sessions			Weekly Spend Given Search Sessions		
	(1)	(2)	(3)	(4)	(5)	(6)
Restaurant Count	−0.001	−0.001	−0.001	0.018	0.010	0.018
	(0.0004)	(0.0004)	(0.0004)	(0.013)	(0.014)	(0.013)
Cuisine Count		0.013**			0.393**	
		(0.005)			(0.144)	
Cuisine Entropy			−0.056			1.650
			(0.089)			(2.830)
Search Indicator Residual	0.104***	0.105***	0.104***	2.876***	2.885***	2.877***
	(0.007)	(0.007)	(0.007)	(0.233)	(0.233)	(0.233)
Count Elasticity	-0.09194	-0.13454	-0.09319	0.11489	0.06536	0.11630
Observations	135,414	135,414	135,414	135,414	135,414	135,414
R^2	0.522	0.522	0.522	0.544	0.544	0.544
Adjusted R^2	0.290	0.290	0.290	0.323	0.323	0.323

Note: *p<0.05; **p<0.01; ***p<0.001
All specs include ZCTA-Week and Individual FEs
Standard Errors clustered at ZCTA-week level

Robustness Checks

Table OA12: Event Study: Impact on Adoption Rate at Census Tract Level

	Dependent variable:	
	Change in Adoption Rate	
	(1)	(2)
Change in Restaurant Count	0.001	0.001
	(0.001)	(0.001)
Change in Cuisine Count		0.001
		(0.001)
Observations	9,309	9,309
R^2	0.157	0.157
Adjusted R^2	0.037	0.037

Note: *p<0.05; **p<0.01; ***p<0.001
All specs include ZCTA-Week and Unit FEs.
Standard Errors are clustered at the ZCTA-Week Level.
Estimated using only restaurant entry
from platform mergers

Table OA13: Event Study: Impact on Churn Rate at Census Tract Level

	Dependent variable:	
	Change in Churn Rate	
	(1)	(2)
Change in Restaurant Count	2.171*	2.197*
	(0.902)	(0.881)
Change in Cuisine Count		2.651
		(3.573)
Change in Promo Usage	11.547	11.601
	(6.696)	(6.701)
ZCTA-Week FE?	N	Y
Observations	3,270	3,270
R^2	0.319	0.319
Adjusted R^2	0.068	0.068

Note: *p<0.05; **p<0.01; ***p<0.001
All Specs include ZCTA-Week FEs
Standard Errors clustered at ZCTA-Week level
Estimated using only restaurant entry from platform mergers

Table OA14: Effect of Assortment Expansion due to Merger on Weekly Orders

	Dependent variable:			
	Weekly Orders		Difference in Weekly Orders	
	(1)	(2)	(3)	(4)
Restaurant Count	−0.0004	−0.0005		
	(0.0003)	(0.0003)		
Cuisine Count		0.0003		
		(0.002)		
Change in Restaurant Count			0.001	0.001
			(0.002)	(0.002)
Lag Restaurant Count			−0.00001	−0.00002
			(0.0001)	(0.0001)
Change in Cuisine Count				−0.002
				(0.006)
Observations	56,483	56,483	56,483	56,483
R^2	0.386	0.386	0.027	0.027
Adjusted R^2	0.208	0.208	−0.005	−0.005

Note: *p<0.05; **p<0.01; ***p<0.001
All Specifications have ZCTA-Week Controls
Standard Errors Clustered at ZCTA-Week Level
Estimated using only restaurant entry
from platform mergers

Table OA15: Effect of Assortment Size on Weekly Orders: Alternative Measures

	Dependent variable:		
	Weekly Orders		
	(1)	(2)	(3)
Restaurant Count	−0.0005***		
	(0.0001)		
Rest. Count at Last Search		−0.00001	
		(0.00004)	
Imputed Rest. Count			−0.0001
			(0.0001)
Weeks Since Last Search	−0.001***	−0.001***	−0.001***
	(0.00003)	(0.00005)	(0.00003)
Observations	1,508,821	1,508,821	1,508,821
R^2	0.299	0.299	0.299
Adjusted R^2	0.264	0.264	0.264

Note: *p<0.05; **p<0.01; ***p<0.001
All specs include Individual and ZCTA-Week FEs
Standard Errors clustered at ZCTA-Week level

Table OA16: Comparison of New and Incumbent Restaurants

Statistic	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)
Panel A: Incumbents						
Price	29.34	8.29	7.99	24.25	27.73	33.13
Delivery Fee	2.43	1.89	0.00	0.77	2.32	3.86
Yelp Rating	3.63	0.43	1.00	3.50	3.50	4.00
Yelp Review Count	776.35	748.32	1	184	403	1,340
Yelp Price Tier	1.63	0.51	1	1	2	2
Sales Quantile	0.66	0.32	0.00	0.49	0.78	0.91
Panel B: Entrants						
Price	26.90	9.42	8.28	20.48	25.70	31.30
Delivery Fee	3.76	1.24	0.00	3.39	3.99	4.45
Yelp Rating	3.52	0.69	1.00	3.50	3.50	4.00
Yelp Review Count	332.63	375.70	0	72	208	467
Yelp Price Tier	1.54	0.57	1	1	2	2
Sales Quantile	0.24	0.29	0.00	0.00	0.00	0.46

OA2 Two Way Fixed Effects Robustness

Recent work has highlighted the possibility of estimation failures in identifying effects for staggered adoption difference-in-differences designs (De Chaisemartin and d'Haultfoeuille, 2020; Sun and Abraham, 2020; Callaway and Sant'Anna, 2020). In particular, guaranteeing estimation of the desired static causal treatment parameter using two-way fixed effects estimation relies on further assumptions around treatment homogeneity, no anticipation of treatment, and treatment dynamics. While numerous estimators have been proposed to address these concerns, none so far can handle the large number of units and the continuous treatment (assortment size) I use in this project. I propose three broad sets of solutions to test the robustness of my estimated effects.

1. Where feasible, test for unobserved heterogeneity in treatment effects (De Chaisemartin and d'Haultfoeuille, 2020)

2. Construct difference-in-difference estimates for each entry experiment without staggered timing

3. Saturate model with observed heterogeneity

In the case of (1), I test the robustness of the adoption specifications (e.g. Table **??**) to unobserved treatment heterogeneity and negative weighting. I find that it is possible that the positive effect of assortment size on adoption is actually negative under sufficient treatment heterogeneity, or that the true effect of assortment size is on average 0 but has positive variance.

The size of my data prohibits me from considering the prior procedure for the effect of additional variety on returning customers. Instead, I will consider (2) and (3) as alternatives. (2) decomposes the continuous treatment of assortment size into its binary parts: specific restaurant entry and exit. I will estimate the marginal effect of each restaurant's entry on returning consumers using only two periods. The entry timing of each restaurant is uniform, so this design will avoid using any staggered treatment timing for each restaurant-specific event study design. However, by construction, I will estimate simple two-period difference-in-difference estimates for the contemporaneous effect of single restaurant entry onto the platform. In particular, I estimate:

$$y_{it} = \alpha_i + \alpha_{z(i)t} + \beta_j D_{ijt} + \epsilon_{it}$$

For restaurant entries j (with treatment dummy D_{ijt}), using only the immediate pre- and postentry weeks for estimation and neighborhood-week controls $\alpha_{z(i)t}$. I expect these estimates to be very noisy. If the true average effect is similar to what I report in the body of the paper, these designs are underpowered (in particular since the treated group is often only several hundred observations. I show the mean and distribution of these effects in Figures OA13 and OA14 for outcome measures weekly orders and weekly spending.

The average effect across restaurants is consistent with the homogeneous effect estimated in the main specifications - restaurant entry reduces the probability of purchase and the average level of spending. However, because the effect varies so much across restaurants, I can't rule out that these effects are consistent with a zero-mean process.

Figure OA13: Restaurant-specific Effect on Orders

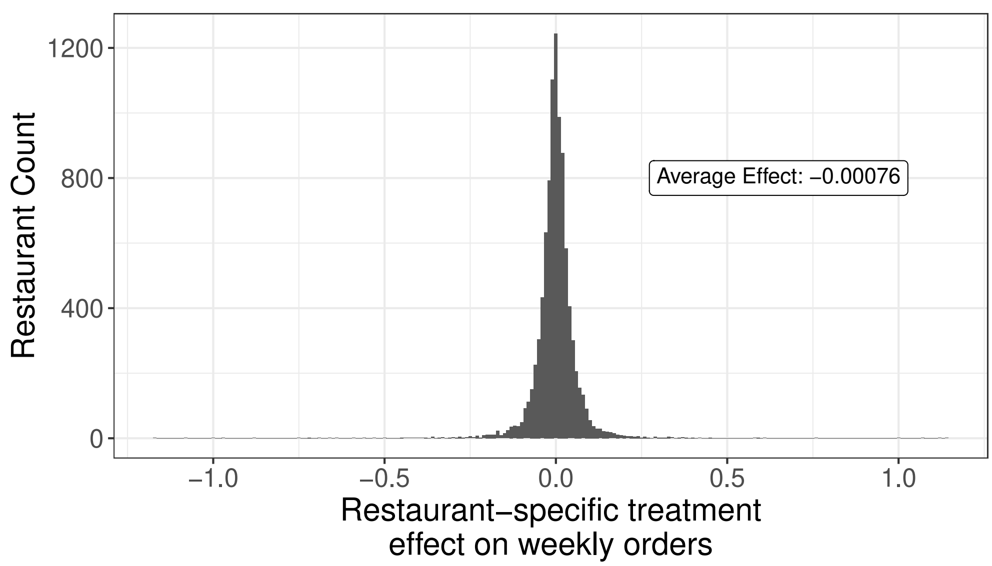

Finally, while I can't directly address the possibility of unobserved treatment heterogeneity, the average effect of assortment size on consumption is mostly consistent (i.e. negative) across many observable sources of heterogeneity. For example, the marginal effect is consistently negative across assortment size, across consumers with different choice histories (excepting consumers who never vary their consumption), and across restaurants of different qualities.

Additionally, I allow for neighborhood-specific effects of assortment size. This may capture the relative pain of choosing online relative to growth of offline choices. Similar to the restaurant-specific effects shown above, the effect at the neighborhood (ZCTA) level is similar on average to the main effects, but it varies widely across neighborhoods. Figure OA15 plots the distribution of effects. This variability is not related on average to the relative number of panelists in the area

Figure OA14: Restaurant-specific Effect on Spending

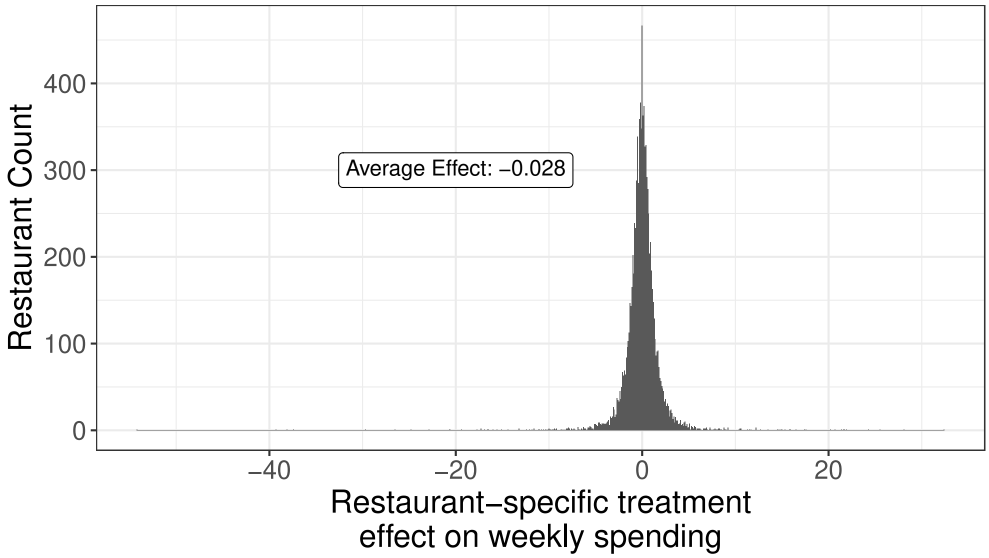

(and thus sample size) or to the average total restaurant online availability in the area.

Figure OA15: Neighborhood-specific Effect on Orders

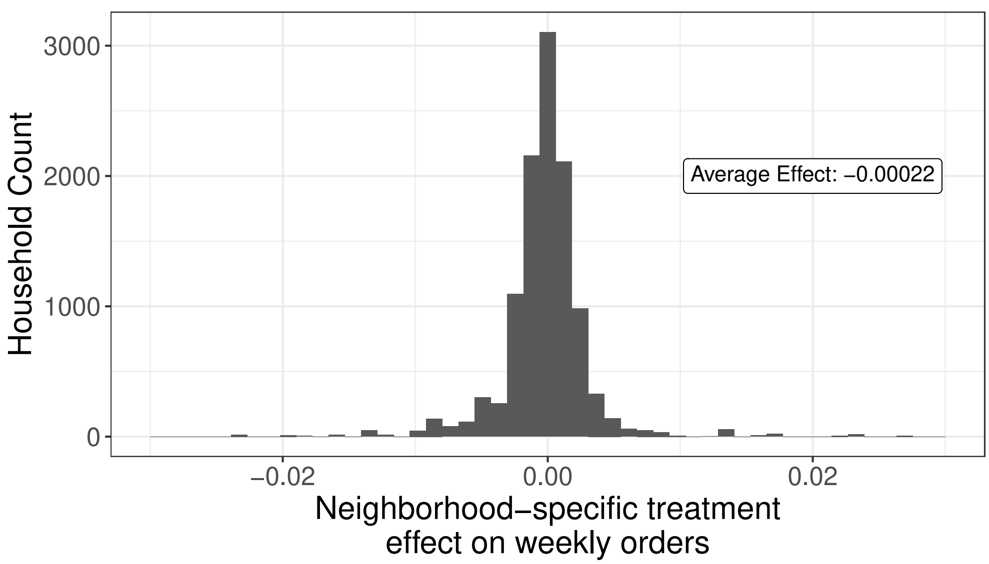

I conclude from these robustness checks that I can rule out large sign reversals from my estimation techniques. However, the average negative impact of assortment size on purchase frequency may be misleading - the two way fixed effects estimation technique may obscure a true effect that is very close to zero.

OA3 Within- and Across-Neighborhood Variation

Much of the variation in assortment sizes is perfectly correlated with the sets of fixed effects I use in my regression analysis. Table OA17 presents the R-squared statistic for a series of fixed effects on two treatment measures: assortment size (number of restaurants) and assortment size changes (number of entering restaurants).

Fixed Effects	Measure	
	Number of Restaurants	Restaurant Entry
Census Tract and Week	0.963	0.353
Census Tract and ZCTA-Week	0.999	0.920
Household and Week	0.943	0.358
Household and ZCTA-Week	0.996	0.912

Table OA17: R-Squared from Fixed Effect Regressions

One main premise of the generalized difference-in-differences identification strategy used here is that consumers who live in the same neighborhood are more comparable than those who live in different neighborhoods. I will show that neighborhoods (measured by ZCTA in this paper), explain a considerable share of variation in consumer demographics.

To test whether census tracts are more similar within or across neighborhoods, I conduct a descriptive exercise by regressing demographic variables on a vector of ZCTA fixed effects. I report in Table OA18 the R-squared and F-statistics from these regressions. ZCTA fixed effects explain a significant and large share of the demographic variation in census tracts.

Table OA18: Variation in Demographics Within ZCTA

	R-Squared	F-Statistic
Pop2015	0.364	4.055
Hh Med Inc	0.672	14.545
Hh Mean Inc	0.691	15.848
Hh 100 To 150	0.501	7.131
Hh 150 To 200	0.562	9.114
Hh Over 200	0.673	14.600
Perc Hs	0.764	22.897
Perc Bach	0.805	29.324
Perc Hs 18to24	0.306	3.124
Perc Bach 18to24	0.532	8.071
Perc Hs Over25	0.649	13.134
Perc Bach Over25	0.771	23.865
Perc Grad Over25	0.769	23.543
Participation Perc	0.449	5.770
Unemployment Rate	0.428	5.314
Pop Black	0.683	15.257
Pop Asian	0.688	15.637
Pop Hispanic	0.643	12.774
Perc Black	0.785	25.820
Perc Asian	0.767	23.355
Perc Hispanic	0.792	26.939

OA4 Selection into Search Data

Only about three quarters of the consumer panel has any match in the search data. I restrict this match further to ensure that the searches accompanying purchase are present at least half of the time (they are otherwise inferred). The subset of users which have search data is not identical to those without - the summary statistics are presented in Table OA19. The means of these summary statistics can reject the null hypothesis that they are the same, but inspection of their magnitudes suggests that the samples are fairly similar.

Table OA19: Comparison of Sample with and without Search

Statistic	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)
Panel A: Search Data						
Orders	18.40	28.24	2	3	8	21
Restaurants Tried	6.91	8.10	1	2	4	9
Spending	487.87	794.61	16.40	81.10	203.86	549.63
Initial Assortment Size	107.09	79.00	1	31	95.5	172
Panel B: No Search Data						
Orders	17.35	27.59	2	3	7	18
Restaurants Tried	6.60	7.98	1	2	4	8
Spending	434.91	744.19	14.70	70.02	159.95	457.88
Initial Assortment Size	97.17	77.57	1	26	68	162

Table OA20: Search Data Summary Statistics

Statistic	Mean	St. Dev.	Median
Weekly Sessions	0.329	1.362	0
Session Duration (minutes)	4.081	4.229	3.000
Search Queries	1.239	0.444	1.000
Desktop Share	0.530	0.451	0.500

Summary statistics of the search data are shown here. Conditional on having interaction with the platform, a user need not enter a search query to end up with a purchase (e.g. navigating from links presented on the home page such as previous purchases). The median user, however, does search via query. Desktop and mobile split total usage evenly in this sample. The typical session duration is quite short - only a few minutes are spent searching.

OA5 Multi-homing and platform entry timing

Many of the restaurants that enter the focal platform in this paper engage in multi-homing: they operate on multiple, competing food delivery platforms. I do not have data from competing platforms, so I cannot measure multi-homing directly. To aid in interpreting the effects I find, I want to rule out that restaurants simultaneously enter multiple platforms—if this were the case, I won't be able to attribute effects solely to changes on the focal platform. Conversations with the platform and with restaurants suggest this is unlikely to occur, but I want to provide some empirical evidence to support this claim.

I collect data on Yelp review text and time stamps to capture the timing and mentions of the focal and competing delivery platforms. In particular, I gather the earliest review date that mentions ordering from the platform. Since the data collection requires some manual review, I sample restaurants for manual review that entered during the sample window before the start of 2018 (allowing for a lag to show up in review data), had matched with a Yelp page, and were not national chains. I sampled 150 restaurants at random from this set and 50 additional randomly sampled restaurants from the top 20% of restaurants on the platform.

I use this data to construct two tests. I normalize the earliest review date for each restaurant-platform combination by converting to the days since the restaurant entered the focal platform. Reviews mentioning the competitor platforms appear on average over a year (404 and 463 days) later than reviews mentioning the focal platform. First, I test whether this time difference is different on average for the focal platform and competing platforms (with a finite sample t-test). I reject that the first reviews occur at the same time for the largest competing platforms. Second, I test whether these differences form different distributions (using a Kolmogorov-Smirnov test). The null hypothesis of no differences across platforms is rejected for the two largest competing platforms, though this test has less power than the simple test of means in the finite data - I can only reject with 95% confidence in one case.

This approach is noisy: only 98 of the 200 restaurants had Yelp review that mentioned the focal platform. In addition, the first occurrence of a review related to receiving delivery may not be always reflecting entry timing in a meaningful way. However, in the absence of additional platforms' data, the review text data are consistent with multi-homing occurring over time, rather than simultaneously when restaurants enter online delivery markets.