HARVARD | BUSINESS | SCHOOL

Shipping Fees and Product Assortment in Online Retail

Chaoqun Chen
Donald Ngwe

Working Paper 19-034

Shipping Fees and Product Assortment in Online Retail

Chaoqun Chen
Southern Methodist University

Donald Ngwe
Harvard Business School

Working Paper 19-034

Shipping Fees and Product Assortment in Online Retail

Chaoqun Chen
Cox School of Business
Southern Methodist University
chaoqunc@smu.edu

Donald Ngwe
Harvard Business School
Harvard University
dngwe@hbs.edu

September 21, 2018

Abstract

Shipping fees are an important aspect of online retail for both consumers and sellers. A common fee structure is *contingent free shipping*, in which consumers are granted free shipping for basket sizes above a minimum value, and are charged a flat fee for orders below this threshold. We seek to characterize how contingent free shipping influences purchase outcomes in a multi-category shopping environment. We build a demand model in which consumers choose how to allocate their spending over different product categories to maximize their direct utility under contingent free shipping. We estimate model parameters using transaction data from a pure online fashion retailer. We find that, relative to offering free shipping, offering contingent free shipping increases basket sizes by encouraging consumers to meet the minimum order threshold. Consumers incur search costs to meet this threshold exactly; sellers may benefit from maintaining high search costs to encourage overshooting. Moreover, we find that contingent free shipping shifts demand to more popular products, and that the effects of category-level price changes on profits depends on the active shipping policy. Our findings demonstrate the importance of jointly determining product assortment attributes and shipping fee policies.

Keywords: shipping fees, e-commerce, multi-category demand, search costs, shipping membership

1 Introduction

The growth of online retail has greatly outpaced that of offline retail in recent years. According to the National Retail Federation, online sales in the United States will grow more than twice as much in 2018 than the retail industry as a whole (Smith, 2018). The shift to e-commerce entails profound changes in the way firms serve their customers. Compared to brick-and-mortar retail, online shopping broadens the assortment available to a customer, enables her to browse a seller's entire product assortment, and allows her to receive her purchase at a location of her choosing.

These differences come hand-in-hand with interesting cost implications. While an online store may offer greater coverage and accessibility for a fraction of what it costs to build and run physical stores, order fulfillment and delivery may cost several times as much for online sellers. In many cases, this may make running physical retail stores more profitable on balance (Reagan, 2017).

Many online sellers seek to recoup delivery costs by charging customers for shipping. Shipping policies come in various forms (Jarvi, 2017). Some websites offer free shipping, either absorbing the delivery cost entirely or raising product prices instead. Others charge flat rates for shipping as a means of encouraging bigger orders. Another widespread tactic is to pass on the exact shipping cost quoted by third-party logistics providers such as UPS or FedEx.

A commonly used combination of these approaches is *contingent free shipping*, in which the seller charges a shipping fee for orders below a minimum order value, but offers free shipping for orders above the minimum. Bloomingdale's, for example, offers free shipping for online orders \$150 and above and charges between \$6 and \$13 for orders below the threshold. Amazon offers a similarly structured shipping fee schedule for its non-Prime customers.[1]

Online sellers balance cost-side considerations with sales implications when choosing how and what to charge for shipping. In one survey, 44% of respondents claimed that their most recent

[1] Amazon Prime is an annual membership sold by Amazon that offers subscribers free two-day shipping for qualifying orders of any amount. This program also has some similarities with contingent free shipping; for example, Amazon designates certain low-value products as "Add-on Items," which must be part of an order of at least \$25 to be eligible for purchase.

reason for abandoning a shopping cart was due to shipping fees that were too high, the most commonly cited reason.[2] On the other hand, offering contingent free shipping may encourage some shoppers to increase their order sizes to meet the minimum order value.

In this paper, we seek to explore the profit implications of different shipping fee policies. We focus on the impact on purchase behavior of the minimum order value and the flat rate shipping fee in a contingent shipping policy regime. We are particularly interested in how shoppers assemble their purchase baskets with products from multiple categories in response to shipping policies, and consequently the interaction of these shipping policies and product assortment in determining overall performance. We focus on environments in which consumers are made aware of the shipping policy upfront, rather than those in which the shipping fee is a "shrouded attribute" (e.g. Brown et al. 2010).

We build a structural demand model to characterize shopper behavior in a setting with a large set of products in multiple categories. We use a direct utility approach, which allows us to jointly characterize a consumer's tradeoffs in allocating her total expenditure among different product categories, any shipping fee payments, and an outside good. We incorporate a novel search cost term, which represents the time and effort it takes for a consumer to meet the minimum order threshold without significantly overshooting.

We estimate demand parameters using a large transaction-level data set from an online fashion retailer in the Philippines. The data set possesses several features that make it highly suitable for addressing our research questions. First, it contains variation in shipping fee policies, including changes in the minimum order threshold, changes in the flat shipping fee, and the occurrence of periods in which shipping is free for all transactions. Second, the data set has a panel structure that allows us to use RFM (recency, frequency, monetary value) variables to account for consumer heterogeneity. Third, the firm has a large product assortment and the data contain detailed category information, which allows us to investigate the relationship between category product assortment

[2] North American Technographics®Retail Online Survey, Q3 2009 (US), Forrester Consulting, accessed at `https://www.ups.com/media/en/Smarter_Strategies_for_Free_Shipping.pdf`

and shipping policies. Finally, the firm is the largest online fashion retailer in the country with no direct competitors, allowing us to abstract from strategic interaction.

Our demand estimates capture several factors affecting online shopping behavior. We detect large differences between product categories in baseline demand and satiation rates. In our setting, for instance, the tops category (e.g. shirts and blouses) has higher baseline demand than the dresses category; however, conditional on purchasing these categories, consumers reach satiation slower when spending on dresses. We also find that shoppers are less sensitive to shipping fees than they are to product prices, a finding that echoes results from the mental accounting literature.

We use our demand estimates to simulate outcomes under alternative shipping policy regimes. We find that contingent free shipping provides strong motivation for consumers to increase their basket sizes. Not only are merchandise profits higher under contingent free shipping than under free shipping, but our estimates imply that the firm could derive even higher merchandise profits by increasing both its flat shipping fee and its minimum order size for free shipping.

In a second set of simulations, we lower the average price of each product category under different shipping policies. Our demand model allows for these price perturbations to affect not only category demand, but also cross-category substitution as well as overall expenditures through income effects. We find that the demand response to these price changes depends heavily on the active shipping policy.

In a third simulation, we explore how manipulating search costs affects purchase behavior under contingent free shipping. Sellers may operationalize such a scenario by, for example, providing product recommendations or adjusting product assortment to make it easier for consumers to meet the free shipping condition without significantly overshooting. Our results suggest that sellers may benefit from maintaining high search costs.

In a final policy simulation, we explore the possibility of offering consumers a membership in which they pay an annual fee in exchange for free shipping on all of their orders for the year. We find that such a program increases the number of orders and the average order size for subscribers.

However, the corresponding loss in shipping revenue may make such a program unprofitable for sellers.

The remainder of the paper proceeds as follows. In Section 2, we review the related literature on shipping fees. In Section 3, we describe the data and empirical setting and provide preliminary evidence for our effects of interest. We specify our demand model in Section 4 and provide estimation details and results in Section 5. We present the results of our policy simulations in Section 6. Section 7 concludes.

2 Related Literature

Shipping policies have been studied from a broad array of consumer behavior and online retail management perspectives. Consequently, the full array of related literature draws from streams in economics, quantitative marketing, consumer psychology, and operations management. In this section, we review the papers and findings from each subfield that are most closely related to our research.

The specific shipping policies most often studied in the literature are unconditional free shipping, contingent free shipping, and graduated shipping fees (in which the shipping fee increases with order size). In some of the earliest work on online shipping fees, Lewis (2006) and Lewis et al. (2006) develop econometric approaches that enable them to analyze the implications of all three types of shipping policies. Of the three policies, they find that an unconditional free shipping policy is the most effective at attracting new customers and increasing order incidence. This finding is further supported by Yao and Zhang (2012), who analyze the joint pricing decision for products and shipping. However, free shipping results in a significant reduction in profits relative to paid shipping models.

Yang et al. (2005) develop a model that verifies the findings of Lewis et al. (2006) and analyzes consumer welfare in response to changes in contingent free shipping thresholds. They find that customers spent an average of \$17 more and purchased 1.82 more items from Amazon under a \$49 contingent free shipping threshold versus a \$25 threshold. The authors also study the implications of price changes under contingent free shipping policies and find that although consumers are faced with greater costs following a price increase, this increases the probability that they will reach the free shipping threshold. In our work, we consider how price effects may differ between categories given contingent free shipping.

Research in consumer psychology finds that consumers place a significant value on products that are free (i.e. the zero-price effect or zero-risk bias); hence, consumers believe they are obtaining a greater value when shipping poses no cost to them (Frischmann et al., 2012; Shampanier et al., 2007). From a retailer perspective, free shipping reduces profits from sales since the retailer is assuming the entire cost of shipping. A subsequent increase in the product price is required to offset this reduction. Frischmann et al. (2012) find that German retailers charge an average of 4.65 euros more under a free shipping policy.

A separate stream of research in consumer behavior has found that a retailer's choice to partition prices has an impact on consumer attentiveness and perception. More specifically, consumers tend to be more sensitive to secondary costs when they are apportioned from the gross price (Bertini and Wathieu, 2008). Displaying shipping charges separately can influence how consumers view a retailer's shipping policy. Consumers become more skeptical of retailers that charge a high shipping fee relative to the product price and believe that these retailers are using their shipping fee as a means of earning additional profit (Chatterjee, 2011).

Research in operations management has accounted for competition in determining optimal shipping policies. Gümüş et al. (2013) examine competition between sellers that offer free shipping and incorporate their shipping costs within product prices and those that charge separate shipping fees. They assume that consumers differ according to how sensitive they are to shipping fees, and characterize optimal seller strategies contingent on the relative size of each group of consumers, the seller's popularity, and its products' cost-to-price ratios.

Leng and Becerril-Arreola (2010) model a two-stage decision process in which the retailer first decides on profit margins and the contingent free shipping threshold, and then determines inventory levels. They find that if retailers decide to use shipping fees as a source of operating revenue, then when a retailer's revenue from shipping decreases it should be met with a reduction in the contingent free shipping threshold and a rise in the product price. If the retailer decides to subsidize a portion of shipping, then the retailer can simultaneously increase the contingent free shipping threshold and its product price. They also find that in a relatively homogeneous market a retailer should set a low contingent free shipping threshold and a high threshold in heterogeneous markets.

In closely related work, Becerril-Arreola et al. (2013) find that a decrease in a retailer's profit margin and the application of a conditional free shipping threshold increases the average order size. Furthermore, they find that as the ratio of inventory to shipping costs decreases it is optimal for the retailer to decrease their profit margin, but at the same time increase the optimal conditional free shipping threshold.

Xu (2016) develops a model designed to examine the implications of various contingent free shipping thresholds. Xu (2016) finds that even though decreasing a contingent free shipping threshold by two-thirds leads to a boost in total sales, it is insufficient in offsetting the overall profit loss, which is about 20 percent. Xu (2016) also finds that decreasing the contingent free shipping threshold does not have a noticeable effect on order counts.

Finally, Jiang et al. (2013) use a non-linear mixed integer model to optimize a retailer's shipping cost decision, taking into account multiple products, customer delivery preferences, and pricing constraints on retailers. Using data on 100 books from Amazon's best-seller list, their model is able to deliver significant consumer savings and retailer profits compared to free shipping and fixed shipping fee strategies.

In our approach, we seek to incorporate much of the insight on shipping policies found in previous research. Specifically, we aim to build a demand model that can characterize the relationship between conversion, order size, and different shipping policies while also allowing for different levels of sensitivity to shipping fees and product prices. We seek to add to the existing literature by demonstrating how shipping fees interact with multi-category product assortment decisions, as well as search costs, to influence purchase outcomes.

3 Data and Preliminary Evidence

We use data provided by an online clothing and accessories retailer in the Philippines. The sample begins at the firm's inception in September 2012 and ends in October 2016. Transaction-level records are available, containing both product and consumer attributes. During this period almost 700,000 customers completed over 2,000,000 orders. Table 1 contains further characteristics of the data set and Table 2 presents summary statistics for completed transactions.

Table 1: Data set characteristics

Sample duration	3 September 2012–31 October 2016
Number of orders	2,116,320
Number of items sold	4,694,023
Number of unique customers	681,318
Number of unique SKUs	972,124

Table 2: Transaction summary statistics

	Mean	**SD**
Order size (PHP)	2,664.06	7,229.87
Selling price	680.15	774.56
Gross margin	143.78	303.18
Discount	17.21%	22.49%

Within the sample duration, the firm makes several changes to its schedule of shipping charges. Table 3 lists each policy change. A key feature of the data is the occurrence of periods in which the firm offers free shipping on all orders. (Appendix Figure A indicates when within our sample these events occur.) The most commonly observed shipping fee schedule charges 100 Philippine pesos (PHP) for orders less than PHP 995, and offers free shipping otherwise. Figure 1 shows the distribution of order sizes under this shipping fee schedule (truncated at PHP 5,000). There is a pronounced bunching of order sizes above the free shipping threshold that does not correspond to the distribution of item prices, which is relatively smooth everywhere.

Table 3: Shipping fee schedules

	Period	**Policy**
A	9/3/2012-10/9/2012	Free shipping on all orders
B	10/9/2012-10/26/2012	Free shipping on orders ≥500 80 on orders <500
C	10/27/2012-12/3/2012	Free shipping on orders ≥1000 100 on orders <1000
D	12/4/2012-4/30/2013	Free shipping on orders ≥1000 100 or 150 on orders <1000 depending on location
E	5/1/2013-11/28/2013	Free shipping on orders ≥1000 100 or 200 on orders <1000 depending on location
F	11/29/2013-9/4/2014	Free shipping on orders ≥1000 100 on orders <1000
G	9/5/2014-10/31/2016	Free shipping on orders ≥995 100 on orders <995
H	Exceptions	Free shipping days

Figure 2 reveals two separate phenomena occurring around the free shipping threshold. This figure overlays the distribution of multi-item purchase baskets over that of single-item purchase baskets. We see that single item purchases spike at the threshold amount, whereas multiple item purchases rise dramatically at the threshold, suggesting that consumers (i) purchase higher-priced items or (ii) purchase additional items in response to the free shipping incentive.[3]

Figure 3 presents further evidence that some consumers "top up" their orders with additional items in order to meet the free shipping threshold. To generate this figure, we limit the data set to orders containing exactly two items. We group these orders in 100-PHP bins and compute the average portion of the basket size accounted for by the less expensive item. We find that the lowest of these averages is found just above the minimum order threshold. This pattern is consistent with

[3]Figure B in the appendix plots the average number of items against order size.

Figure 1: Histogram of order sizes

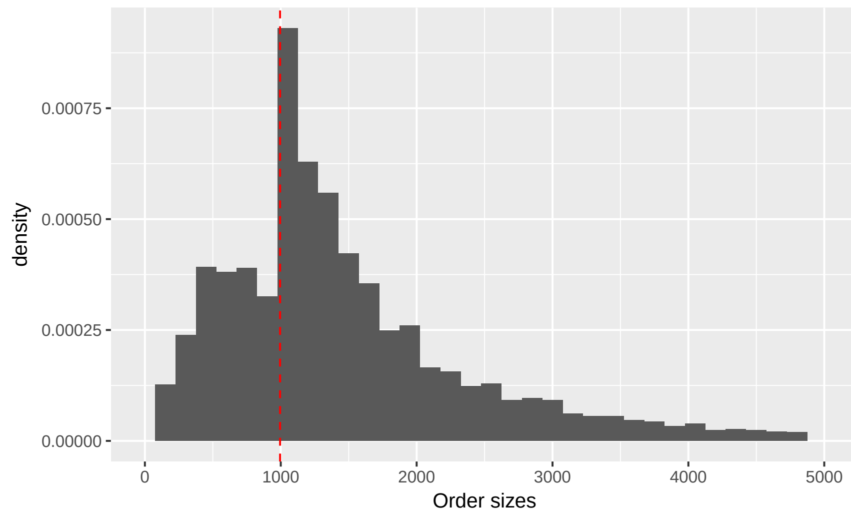

Figure 2: Distribution of single- and multiple-item orders

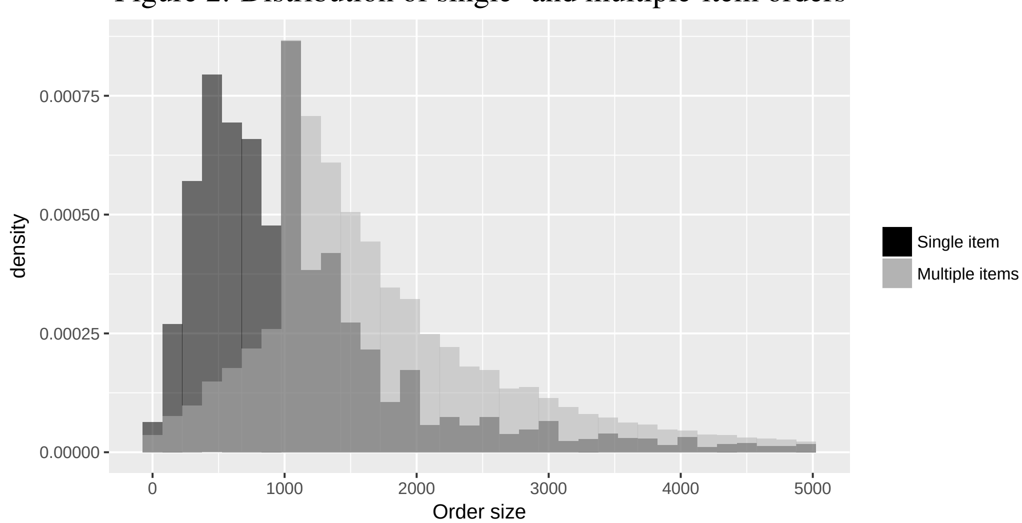

Figure 3: Average portion of order of less expensive product within two-item orders

consumers purchasing lower-priced items partly to avoid paying the flat rate shipping fee.

The firm sells items from nine mutually exclusive product categories. These are listed in Table 4, together with average category attributes. Purchase tendency is the percentage of transactions that contain at least one item in that category. We compute the number of items available as the number of unique SKUs in each category observed monthly in our transaction data. Store brand is the percentage of these items under the firm's private label. Finally, discount frequency is the average percent of products that are discounted at any given period. Appendix Figure C illustrates variation over time in prices.

Table 4: Category Differences

Category	**Purchase tendency**	**Price (PHP)**	**Margin**	**No. items (in 000's)**	**Store brand**	**Discount frequency**
Accessories	0.138	663.011	0.213	5.896	0.050	0.585
Bags	0.081	867.620	0.171	2.615	0.200	0.572
Beauty	0.066	402.149	0.170	2.400	0.001	0.219
Bottoms	0.118	749.613	0.175	6.247	0.182	0.513
Dresses	0.202	862.309	0.235	8.408	0.373	0.537
Shoes	0.406	1,015.422	0.206	21.498	0.232	0.625
Tops	0.285	592.900	0.189	19.991	0.202	0.479
Underwear	0.025	403.393	0.188	2.667	0.015	0.283
Other	0.088	616.514	0.199	11.612	0.038	0.430

Figure 4: Category unit sales shares by expenditure level

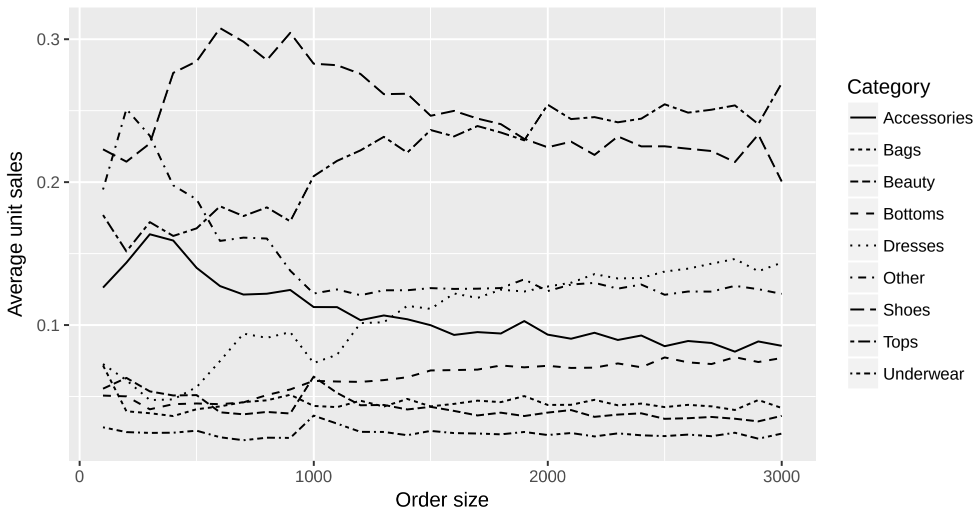

We attempt to find direct evidence in the data for differences in how likely products in each category might be bought to top up purchase baskets to meet the free shipping threshold. Figure 4 shows how consumers allocate their purchase decisions according to categories depending on the spending level. The tops category experiences a bump at the standard minimum order threshold in unit shares, suggesting that consumers tend to focus on this category when they top up their baskets. Similarly, the beauty and underwear categories, both of which tend to be lower-priced categories, show a pronounced spike in their proportion of units sold at the free shipping threshold. This jibes with the pattern in Figure 3.

While these data trends strongly suggest specific purchase patterns in response to contingent shipping fees, descriptive statistics or reduced-form approaches alone are insufficient for characterizing the full extent of consumer behavior in this context. Orders exceeding the minimum order threshold, for instance, may either result from topping up or occur incidentally given the underlying demand for products. Likewise, price variations between product categories may jointly influence the likelihood of topping up. In the following section, we specify a structural demand model in order to account for these and other relationships in the data.

4 Model

In this section, we introduce a model of demand for an online multi-category retail environment. Consumers make decisions sequentially according to a two-stage budgeting process. In each period, a consumer decides on her total expenditure and how to allocate it across categories. Under a contingent free shipping policy, consumers have an incentive to meet the minimum order requirement in order to avoid the shipping fee; however, meeting the threshold without substantially overshooting comes with a search cost. Consumers consider this search cost, category attributes, the shipping fee, and product prices in making final purchase decisions.

4.1 Utility

In each visit to the website, consumer i chooses total spending y and category shares $s_1, \ldots,$ $s_c, \ldots, s_C$[4] to maximize her direct utility:

$$\max_{y,\boldsymbol{s}} E_\epsilon \left[U^C(\boldsymbol{s}, y, \boldsymbol{X}_2, \boldsymbol{p}, \boldsymbol{\epsilon}; \boldsymbol{\theta_2})\right] + \psi_0 \log\left(I - y - C_{ship}(y)\right) - C_{search}(y)$$

$$\text{subject to.} \quad y \geq 0, I - y - C_{ship}(y) \geq 0, s_c \geq 0, \text{for } c = 1, \ldots, C$$

where $U^C(\boldsymbol{s}, y, \boldsymbol{X}_2, \boldsymbol{p}, \boldsymbol{\epsilon}; \boldsymbol{\theta}_2)$ is the utility she derives from allocating y into product category shares, which is a function of category attributes $\boldsymbol{X}_2 = \{X_{21}, \ldots, X_{2C}\}$, average category prices $\boldsymbol{p} = \{p_1, \ldots, p_C\}$, random shocks associated with each category $\boldsymbol{\epsilon} = \{\epsilon_1, \ldots, \epsilon_C\}$, and parameter vector $\boldsymbol{\theta}_2$.[5] The second term of the utility function represents the utility from spending on the outside good, $I - y - C_{ship}(y)$, where $C_{ship}(y)$ is what the consumer pays for shipping.[6] The third term captures the search cost that consumers incur while shopping, which is higher for consumers who seek to reach the minimum order requirement without overshooting.

[4] For simplicity, we drop subscript i.
[5] Product attributes and prices are monthly sales-weighted averages.
[6] We assume an annual income of PHP 250,000, which is roughly the national average.

We assume that total expenditure decisions y are made prior to the realization of category-specific random terms $\boldsymbol{\epsilon}$. Consequently, in our model the random terms $\boldsymbol{\epsilon}$ affect only decisions on expenditure allocation but not the total expenditure. Under this assumption, demand can be represented as a two-stage decision process: consumers first decide on the total spending for a given order, and then decide how much to spend on each category. Specifically, the consumer's problem is:

Stage 1: Total expenditure

$$
\begin{aligned}
&\max_{y} \omega\left(y, \boldsymbol{X}_2, \boldsymbol{p}; \boldsymbol{\theta_2}\right) + \psi_0 \log\left(I - y - C_{ship}(y)\right) - C_{search}(y) \\
&\text{subject to.} \quad y \geq 0, I - y - C_{ship}(y) > 0
\end{aligned}
\tag{1}
$$

where the expenditure maximized utility from allocation, or inclusive value function, is written as $\omega\left(y, \boldsymbol{X}_2, \boldsymbol{p}, \boldsymbol{\epsilon}; \boldsymbol{\theta_2}\right)$, and

Stage 2: Category allocation

$$
\begin{aligned}
&\max_{\boldsymbol{s}} \sum_{c=1}^{C} \psi_c \gamma_c \log\left(\frac{s_c y}{p_c \gamma_c} + 1\right) \\
&\text{subject to.} \quad s_c \geq 0, c = 1, \ldots, C, \sum_{c=1}^{C} s_c = 1
\end{aligned}
\tag{2}
$$

4.2 Specifications

In Stage 1, we specify preference for the outside good as $\psi_0 = \exp(X_0\alpha_0 + \epsilon_0)$, where X_0 are consumer attributes and period variables, and assume that ϵ_0 follows a type I extreme value distribution. During free shipping periods, $C_{ship}(y) = 0$ for all orders; under a contingent free shipping policy with flat rate c_f, shipping cost is a step function:

$$C_{ship}(y) = \begin{cases} \psi_1 c_f, & y < MO \\ 0, & y \geq MO \end{cases}$$

We specify the shipping fee sensitivity such that $\psi_1 = \exp(X_1\alpha_1)$, where X_1 are consumer attributes. The value of ψ_1 governs the consumer's sensitivity to shipping fees. If $\psi_1 = 1$, consumers internalize shipping fees into their total spending equivalently to product prices; on the other hand, $\psi_1 > 1$ or $\psi_1 < 1$ imply that consumers either overrate or underrate shipping fees relative to product prices.

Consumers incur search costs during the shopping process. Moreover, they may experience additional search costs in the process of topping up their purchase baskets to meet the minimum order requirement. Therefore, we specify search cost as follows:

$$C_{search}(y) = \begin{cases} \nu_0, & y = 0 \\ \nu_1, & 0 < y < MO \\ \frac{\psi_2}{y - MO + 1} + \nu_2, & y \geq MO \end{cases}$$

where ν_j, $j = 1, 2, 3$ are random terms that follow a type I extreme value distribution with scale parameter σ_ν. For $y \geq MO$, the search cost is higher for orders that come closer to the minimum order requirement. We specify $\psi_2 = \exp(X_2\alpha_2)$, where X_2 are category-specific attributes. We denote the vector of Stage 1 parameters as $\boldsymbol{\theta}_1 = \{\alpha_0, \alpha_1, \alpha_2\}$.

In Stage 2, the utility function follows the Multiple Discrete-Continuous Extreme Value (MDCEV) framework introduced by Bhat (2008). This utility function assumes that the utility from each category is separately additive (e.g. Shriver and Bollinger 2015). The quantity index for category c, $\frac{s_c y}{p_c}$, scaled by satiation parameter γ_c, enters utility logarithmically. The baseline parameter, ψ_c, can be interpreted as the marginal utility at $s_c = 0$. The greater ψ_c is, the more likely it is that a consumer will buy category c, and we specify it as a function of category attributes:

$$\log(\psi_c) = X_{2c}\beta + \epsilon_c$$

where each ϵ_c is assumed to independently follow a type I extreme value distribution with scale parameter σ_ϵ. Satiation parameters, $\boldsymbol{\gamma} = \{\gamma_1, \ldots, \gamma_c\}$, govern the curvature of the utility function, with a smaller value of γ meaning greater satiation. We require both $\boldsymbol{\psi}$ and $\boldsymbol{\gamma}$ to be positive. We denote the vector of Stage 2 parameters as $\boldsymbol{\theta}_2 = \{\beta, \gamma\}$. For identification, we fix the intercept of one category to be 0.

4.3 Estimation

A key challenge in estimating $\boldsymbol{\theta}_1$ is that the first-stage expenditure decision requires knowing the inclusive value function $\omega(y, \boldsymbol{X}_2, \boldsymbol{p}; \boldsymbol{\theta}_2)$. A joint estimation of the multi-stage model involves a nested procedure in which we need to solve for the inclusive value function for each parameter update. This is highly computationally expensive. To address this problem, we adopt a three-step estimation procedure similar to Hendel and Nevo (2006):

Step 1: We estimate $\boldsymbol{\theta}_2$ in the allocation stage conditional on realized total expenditures.
Step 2: We use $\boldsymbol{\theta}_2$ to simulate and model the inclusive value function.
Step 3: Given the inclusive value function, we estimate $\boldsymbol{\theta}_1$ in the total expenditure stage.

Likelihood for the Stage 2 model To solve for the optimal allocation of the constrained optimization problem in Equation (2), we can write the Lagrangian equation accounting for the constraints:

$$L = \sum_{c=1}^{C} \psi_c \gamma_c \log\left(\frac{s_c y}{p_c \gamma_c} + 1\right) + \mu(1 - \sum_{c=1}^{C} s_c)$$

where μ is the Lagrangian multiplier. The Kuhn-Tucker conditions imply that at the optimal allocation $\boldsymbol{s}^* = \{s_1^*, \ldots, s_C^*\}$,

$$\frac{\partial L(s_c^*)}{\partial s_c} \leq 0$$

The equality holds only when $s_c^* > 0$. Let c' be one category with positive expenditure. Solving this condition yields the following necessary conditions:

$$\begin{cases} \epsilon_c - \epsilon_{c'} < \tilde{V}_{c'} - \tilde{V}_c, & \text{and } s_c = 0 \\ \epsilon_c - \epsilon_{c'} = \tilde{V}_{c'} - \tilde{V}_c, & \text{and } s_c > 0 \end{cases}$$

where

$$\tilde{V}_c = X_{2c}\beta - \log\left(\frac{s_c}{\gamma_c} + \frac{p_c}{y}\right)$$

Consumers may only purchase a subset of categories, leaving the likelihood of $\boldsymbol{s}$ to be a mixture of discrete and continuous variables. For $s_c > 0$, the distribution has a well-defined continuous density function; when $s_c = 0$, it leaves a mass point in the distribution. The probability of one observation can be written as:

$$\begin{aligned} P(s_1, \ldots, s_C) = & \prod_{c:s_c>0} \text{density}(s_c) * \prod_{c:s_c=0} \text{probability}(s_c = 0) \\ = & \frac{1}{\sigma_\epsilon^{M-1}} \left(\sum_{c:s_c>0} s_c + \frac{\gamma_c p_c}{y}\right) \left(\prod_{c:s_c>0} \frac{1}{s_c + \frac{\gamma_c p_c}{y}}\right) \\ & \left[\frac{\prod_{c:s_c>0} e^{\tilde{V}_c/\sigma_\epsilon}}{\left(\sum_{c=1}^{R} e^{\tilde{V}_c/\sigma_\epsilon}\right)^M}\right] (M-1)! \end{aligned}$$

where $M = \sum_c^C 1 * (s_c > 0)$. The log-likelihood function of all the expenditure share across consumers and over time is thus:

$$\begin{aligned}\mathbb{L}_2 =& \sum_i^I \sum_{t=1}^{T_i} ll_{it} \\ \text{with } ll_{it} =& (1 - M_{it}) \log(\sigma_\epsilon) + \log\left(\sum_{r:s_{itc}>0} s_{itc} + \frac{\gamma_r p_{rt}}{y_{it}}\right) + \sum_{r:s_{itc}>0} \frac{1}{s_{itc} + \frac{\gamma_r p_{rt}}{y_{it}}} + \\ & \sum_{r:s_{itc}>0} \tilde{V}_{itc}/\sigma_\epsilon - M_{it}\left(\sum_{r=1}^R e^{\tilde{V}_{itc}/\sigma_\epsilon}\right) + \log\left((M_{it} - 1)!\right)\end{aligned}$$

We use a standard numerical optimization procedure to find $\hat{\boldsymbol{\theta}}_2$ that maximizes $\mathbb{L}_2$.

Modeling the inclusive value function We use our estimates of the structural parameters in the second-stage model $\boldsymbol{\theta}_2 = (\boldsymbol{\beta}, \boldsymbol{\gamma})$, to compute the inclusive value function. As there is no closed form solution for the inclusive value function, we overcome this challenge by first simulating optimal shares conditional on expenditure levels and category attributes $\boldsymbol{X}_2$, then approximating the inclusive value function as a function of expenditure y and category attributes $\boldsymbol{X}_2$. The procedure is detailed as follows.

We follow Pinjari and Bhat (2010) to solve for the optimal shares at a given expenditure level. Then we compute inclusive value with:

$$\begin{aligned}\omega(y, \boldsymbol{X}_2, \boldsymbol{p}, \boldsymbol{\theta}_2) &= \mathbb{E}_{\boldsymbol{\epsilon}}\left[\max_{\boldsymbol{s}} U^C(\boldsymbol{s}; y, \boldsymbol{X}_2, \boldsymbol{p}, \boldsymbol{\epsilon}, \boldsymbol{\theta}_2)\right] \\ &= \int_{\boldsymbol{\epsilon}} U^C(\boldsymbol{s}^*; y, \boldsymbol{X}_2, \boldsymbol{p}, \boldsymbol{\epsilon}, \boldsymbol{\theta}_2) dF(\boldsymbol{\epsilon})\end{aligned}$$

where $\boldsymbol{s}^*$ is the optimal vector of category shares given expenditure level y, category attributes $\boldsymbol{X}_2$ and price $\boldsymbol{p}$, as well as $\boldsymbol{\epsilon}$. We obtain the integration with sparse grid quadrature, that is,

$$\omega(y, \boldsymbol{X}_2, \boldsymbol{p}, \boldsymbol{\theta}_2) = \sum_k w_k U^C(\boldsymbol{s}^*; y, \boldsymbol{X}_2, \boldsymbol{p}, \boldsymbol{\epsilon}^{(k)}, \boldsymbol{\theta}_2)$$

where $\boldsymbol{\epsilon}^{(k)}$ is the k-th node of a sparse grid (Heiss and Winschel, 2008), and w_k is the weight associated with k-th nodes.

Another challenge arises due to the large dimensionality in $\boldsymbol{X}_2$ and $\boldsymbol{p}$. It is very computationally expensive to obtain the inclusive value function at all possible combinations of expenditure, category attributes, and price; we alleviate this computational burden by Chebyshev polynomial interpolation.

Likelihood for the Stage 1 model Due to the nature of the step function in the first-stage model, the optimal expenditure could occur in three regions: $y^*_{(0)} = 0$, $0 < y^*_{(1)} < MO$, or $y^*_{(2)} \geq MO$. The optimal expenditure needs to satisfy the first order condition, $\frac{\partial U(y^*)}{\partial y} = 0$, for $y^* > 0$ in the transaction data. With our empirical specification, that is,

$$\omega_y(y^*_{(1)}, \boldsymbol{X}_2, \boldsymbol{p}, \boldsymbol{\theta}_2) - \frac{\psi_0}{I - y^*_{(1)} - \psi_1 c_f} = 0, \text{ for } 0 < y^*_{(1)} < MO \tag{3}$$

$$\omega_y(y^*_{(2)}, \boldsymbol{X}_2, \boldsymbol{p}, \boldsymbol{\theta}_2) - \frac{\psi_0}{I - y^*_{(2)}} + \frac{\psi_2}{(y - MO + 1)^2} = 0, \text{ for } y^*_{(2)} \geq MO \tag{4}$$

Rearranging the terms yield $\epsilon_0 = h(y)$, and

$$h(y^*) = \begin{cases} -\psi_0 + \log\left[\omega_y(y^*_{(1)}, \boldsymbol{X}_2, \boldsymbol{p}, \boldsymbol{\theta}_2)(I - y^*_{(1)} - \psi_1 c_f\right], \text{ for } 0 < y^*_{(1)} < MO \\ -\psi_0 + \log\left[\left(\omega_y(y^*_{(2)}, \boldsymbol{X}_2, \boldsymbol{p}, \boldsymbol{\theta}_2) + \frac{\psi_2}{(y - MO + 1)^2}\right)(I - y^*_{(2)})\right], \text{ for } y^*_{(2)} \geq MO \end{cases}$$

Meanwhile, the optimal solution also needs to satisfy $U(y^*) = \max_j U(y^*_{(j)})$. To be more specific, for $y^* = y^*_{(j)}$, we have $U(y^*_{(j)}) \geq U(y^*_{(k)}), k \neq j$, and that is,

$$\begin{aligned} &U(y^*_{(j)}) = V_j + \nu_j \geq U(y^*_{(k)}) = V_k + \nu_k \\ &\text{or } \nu_j - \nu_k \geq V_k - V_j \end{aligned}$$

with $V_j = \omega\left(y^*_{(j)}, \boldsymbol{X}_2, \boldsymbol{p}; \boldsymbol{\theta_2}\right) + e^{X_0\alpha_0 + h(y^*)} \log\left(I - y - C_{ship}(y^*_{(j)})\right) - C_{search}(y^*_{(j)})$. Note that while we know the observed expenditure level $y^* = y^*_{(j)}$, we don't know the counterfactual $y^*_{(k)}$. We solve for counterfactual $y^*_{(k)}$ by finding the roots of Equation (4).

Hence, the density of optimal expenditure is:

$$
\begin{aligned}
f(y) =& P(y^*=0)^{d_0} * \left[P(0<y^* \le MO) f_1(y^*_{(1)})\right]^{d_1} * \left[P(y^* \ge MO) f_2(y^*_{(2)})\right]^{d_2} \\
=& \left[\frac{\exp(\frac{V_0}{\sigma_\nu})}{\sum_j \exp(\frac{V_j}{\sigma_\nu})}\right]^{d_0} * \left[\frac{\exp(\frac{V_1}{\sigma_\nu})}{\sum_j \exp(\frac{V_j}{\sigma_\nu})} |J_1| \lambda\left(\frac{h(y)}{\sigma_{\epsilon_0}}\right)\right]^{d_1} * \left[\frac{\exp(\frac{V_2}{\sigma_\nu})}{\sum_j \exp(\frac{V_j}{\sigma_\nu})} |J_2| \lambda\left(\frac{h(y)}{\sigma_{\epsilon_0}}\right)\right]^{d_3}
\end{aligned}
$$

where $d_j = 1 * (y^* = y^*_{(j)})$, and $\lambda(\cdot)$ is the density function of the type I extreme value distribution. J_1 and J_2 are Jacobian matrices, particularly, $J_1 = \frac{\partial \epsilon_0}{\partial y^*_{(1)}} = \frac{\omega_{yy}}{\omega_y} - \frac{1}{I - y - \psi_1 \text{fee}}$, and $J_2 = \frac{\partial \epsilon_0}{\partial y^*_{(2)}} = \frac{\omega_{yy} - \frac{2\psi_2}{(y-MO+1)^3}}{\omega_y + \frac{\psi_2}{(y-MO+1)^2}} - \frac{1}{I-y}$.

Hence, the log likelihood under the contingent shipping policy is:

$$
ll = \sum_j d_j \frac{V_j}{\sigma_\nu} - \log\left(\sum_j e^{\frac{V_j}{\sigma_\nu}}\right) + \sum_{j=1}^{2} d_j * \left[\log(|J_j|) - \log(\sigma_{\epsilon_0}) - \frac{h(y)}{\sigma_{\epsilon_0}} - e^{-\frac{h(y)}{\sigma_{\epsilon_0}}}\right]
$$

When free shipping is offered, the log likelihood is expressed as:

$$
ll^{free} = \sum_{j=0}^{1} d_j \frac{V_j}{\sigma_\nu} - \log\left(\sum_j e^{\frac{V_j}{\sigma_\nu}}\right) + d_1 * \left[\log(|J|) - \log(\sigma_{\epsilon_0}) - \frac{h(y)}{\sigma_{\epsilon_0}} - e^{-\frac{h(y)}{\sigma_{\epsilon_0}}}\right]
$$

where $j = 0$ represents no orders and $j = 1$ denotes an order, $h(y^*) = -\alpha_0 + \log(\omega_y(y^*, \boldsymbol{X}_2, \boldsymbol{p}, \boldsymbol{\theta}_2)) + \log(I - y^*)$, and the Jacobian is $|J| = \frac{\omega_{yy}}{\omega_y} - \frac{1}{I-y}$. $V_0 = \omega\left(0, \boldsymbol{X}_2, \boldsymbol{p}; \boldsymbol{\theta_2}\right) + e^{X_0\alpha_0 + h(y^*)} \log\left(I\right)$ and $V_1 = \omega\left(y^*, \boldsymbol{X}_2, \boldsymbol{p}; \boldsymbol{\theta_2}\right) + e^{X_0\alpha_0 + h(y^*)} \log\left(I - y^*\right)$.

Incorporating website traffic data We observe website traffic data $M_1, \ldots, M_t, \ldots, M_T$, where M_t is the number of user visits to the website during period t. Out of M_t website visits, we observe n_t transactions in the sales data; therefore, in each period t, we count $M_t - n_t$ no-purchase observations.

Our model predicts a no-purchase observation when

$$\begin{aligned} & \omega_y < \frac{\psi_0}{I} \\ \text{or} \qquad & \omega_y I < \exp(X_0\alpha_0 + \epsilon_0) \end{aligned}$$

Therefore,

$$\begin{aligned} P(y=0) &= P(\epsilon_0 > \log(\omega_y I) - X_0\alpha_0) \\ &= 1 - \exp\left[-\exp\left(-\frac{\log(\omega_y I) - X_0\alpha_0}{\sigma_{\epsilon_0}}\right)\right] \end{aligned}$$

The full likelihood for Stage 1 is therefore

$$\begin{aligned} \mathbb{L}_1 = & \sum_{i=1}^{I}\sum_{t=1}^{T_i}\left[(1-Free_t)ll_{it} + Free_t * ll_{it}^{free}\right] \\ & + \sum_{t=1}^{T}(M_t - n_t)\log\left\{1-\exp\left[-\exp\left(-\frac{\log(\omega_{y,t} I) - X_{0t}\alpha_0}{\sigma_{\epsilon_0}}\right)\right]\right\} \end{aligned} \tag{5}$$

where $Free_t = 1$ during free shipping days and $Free_t = 0$ otherwise.

Reweighting As in many online settings, conversion rates are quite low when we take website traffic to correspond to market size. King and Zeng (2001) discuss how logit-type models can sharply underestimate the probability of rare events. We adjust our estimation model in line with their proposed solution. First, we form a subset of the data that includes all the transactions and a random sample of no-purchase observations. We then reweight the log-likelihood function as follows:

$$\mathbb{L}_1' = \sum_{i=1}^{I}\sum_{t=1}^{T_i} w_1 * \left[(1-Free_t)ll_{it} + Free_t * ll_{it}^{free}\right] + \sum_{t=1}^{T} w_0 * (M_t - n_t)\log\left\{1-\exp\left[-\exp\left(-\frac{\log(\omega_{y,t}I) - X_{0t}\alpha_0}{\sigma_{\epsilon_0}}\right)\right]\right\} \tag{6}$$

where $w_1 = \frac{\tau}{\tau'}$ and $w_0 = \frac{1-\tau}{1-\tau'}$, with τ and τ' being the conversion rates before and after resampling, respectively.[7] We use a standard numerical optimization procedure to find $\hat{\boldsymbol{\theta}}_1$ that maximizes $\mathbb{L}_1'$.

5 Demand estimates

In this section, we present and discuss our demand estimates. We begin by discussing the estimated coefficients relating to the category allegation stage of consumer choice, and continue to discuss the estimated coefficients that affect the total expenditure decision. We find substantial heterogeneity in the attractiveness of each product category, both in terms of how likely each category is included in the consumer's purchase basket as well as in how much is spent on each category. We find, as existing research suggests, that consumers do not fully internalize shipping fees. Moreover, we find support for the conjecture that consumers face higher than usual search costs when aiming to reach the minimum order threshold, and that these costs are a function of product assortment.

Stage 2 estimates Table 5 presents the coefficient estimates of the Stage 2 model, which describes expenditure allocation into category shares. Category intercepts, listed first in column β, reflect each category's baseline attractiveness to consumers. Shoes, the category with the highest average budget share, has the highest estimated intercept. Underwear, on the other hand, is the least bought category, and has the lowest intercept.

The bottom three rows of Table 5 contain coefficient estimates for category attributes. We find

[7] The conversion rate before resampling is 2.6%, we sample no-purchase observations such that $\tau' = 0.5$.

Table 5: Coefficient Estimates of the Stage 2 Model

	β	γ
Accessories	0.173^{***}	26.008^{***}
	(0.002)	(0.008)
Bags	-0.256^{***}	32.958^{***}
	(0.003)	(0.009)
Beauty	-1.335^{***}	27.181^{***}
	(0.004)	(0.036)
Bottoms	-0.214^{***}	21.281^{***}
	(0.003)	(0.029)
Dresses	0.304^{***}	47.863^{***}
	(0.003)	(0.055)
Shoes	1.435^{***}	80.906^{***}
	(0.003)	(0.067)
Tops	0.605^{***}	29.904^{***}
	(0.003)	(0.024)
Underwear	-2.098^{***}	22.529^{***}
	(0.006)	(0.055)
Other	0.000	26.797^{***}
	(0.000)	(0.029)
No. of items	-0.003^{***}	
	(0.0002)	
Store brand	1.527^{***}	
	(0.008)	
Discount	-0.049^{***}	
	(0.009)	
Note:	*p<0.1; **p<0.05; ***p<0.01	

that the number of product items within a category does not necessarily lead to higher purchase tenancy, and the estimate is slightly negative. We note that, as in many online retail environments, a large subset of products are rarely bought, which jibes with the "long tail" or "fat tail" effect identified in the literature (e.g., Brynjolfsson et al. 2011). Unlike the number of product items, the fraction of store brand items has a large positive effect on category utility. This result suggests that the retailer's store brands boost category-level sales. Meanwhile, discounting has a negative effect on category attractiveness, controlling for average selling prices. This likely reflects product assortment effects, given that the firm tends to offer discounts on older and weaker-selling items.

The second column of Table 5 presents the estimates of satiation parameters γ. Recall that the larger the satiation parameter, the lower the effect of satiation of consumption in a category, and hence the larger the expenditure on the category conditional on purchase incidence. The estimates are thus consistent with patterns in the sales data, which show that shoes, dresses, and bags are the three categories with the highest average purchase amounts conditional on purchase incidence.

Stage 1 estimates Table 6 presents parameter estimates from the Stage 1 model, which describes how consumers determine their total expenditure on the website. The first column contains coefficient estimates for components of α_0, the deterministic part of ψ_0, which scales the utility of the outside good. Higher values of ψ_0 correspond to lower demand for the retailer's goods. We include July and December dummies to control for seasonal demand peaks and troughs relevant to the local retail market, as well as a quadratic time trend. The coefficient estimates for RFM variables suggest stickiness in purchase behavior: in any given period, a consumer is more likely to spend more on the website the higher the size of her previous purchase basket (monetary value), the longer it has been since her last purchase (recency), and the more frequently she has visited the store in the past (frequency).

The estimated coefficients for components of ψ_1 capture consumer sensitivity to shipping fees. They indicate that more valuable (in the sense of RFM) consumers not only have higher demand, but are also less sensitive to shipping fees. Given our estimates, the average value of ψ_1 over

Table 6: Coefficient Estimates of the Stage 1 Model

α_0		α_1		α_2	
Intercept	12.304^{***}	**Intercept**	-5.497^{***}	**Intercept**	2.377^{***}
	(0.040)		(0.739)		(0.026)
Period	0.340^{***}	**Money**	-0.305	**No. items of Accessories**	-0.117
	(0.040)		(0.914)		(0.151)
Period^2	-0.420^{***}	**Recency**	-0.608^{***}	**No. items of Bags**	-0.013
	(0.041)		(0.123)		(0.189)
July	0.023^{***}	**Frequency**	-0.142^{**}	**No. items of Beauty**	-0.702^{***}
	(0.003)		(0.071)		(0.154)
December	-0.033^{***}			**No. items of Bottoms**	-0.274
	(0.004)				(0.316)
Money	-0.628^{***}			**No. items of Dresses**	0.572^{***}
	(0.017)				(0.164)
Recency	-0.175^{***}			**No. items of Shoes**	0.220^{*}
	(0.012)				(0.113)
Frequency	-0.062^{***}			**No. items of Tops**	-0.557
	(0.017)				(0.371)
				No. items of Underwear	1.318^{***}
					(0.268)
				No. items of Other	-0.680^{***}
					(0.125)
$\log(\sigma_\epsilon)$	0.519^{***}				
	(0.019)				
$\log(\sigma_\nu)$	2.286^{***}				
	(0.054)				

Note: *p<0.1; **p<0.05; ***p<0.01

consumers and periods is 0.1. This is consistent with previous research that finds that consumers tend to be less sensitive to the cost of shipping relative to that of goods (e.g., Morwitz et al. 1998; Brown et al. 2010; Einav et al. 2015; Hossain and Morgan 2006; Tyan 2005).

We also find support for higher than usual search costs near the free shipping threshold. We estimate coefficients for the number of items in each category to test the conjecture that if consumers have more options to choose from, it might be easier for them to fill the shopping cart at close to the exact minimum order requirement. We find that, while this might be true for some categories, search costs seem to be positively correlated with the number of items in other categories. A possible reason for this may lie in the relative price levels and variation between categories; shoes in the website's catalogue, for instance, are priced mostly above the free shipping threshold.

6 Counterfactual simulations

In this section, we simulate a set of counterfactual scenarios with the aim of shedding light on the relationship between shipping fees, product assortment, and demand outcomes. We begin by comparing sales outcomes between the free shipping case and different contingent free shipping schedules. Next, we show how category-level price changes interact with shipping policies to influence sales. We then illustrate how influencing shopper search costs can affect different performance outcomes for the seller. Finally, we simulate a free shipping subscription model.

6.1 Contingent free shipping and demand outcomes

Relative to free shipping, contingent free shipping may encourage some consumers to increase their basket sizes, but cause other consumers to opt for no purchase (Lewis et al. 2006). To quantify this tradeoff for our empirical context, we simulate consumer orders under three shipping policies: (1) free shipping, (2) the "baseline" contingent free shipping schedule most often observed in our data, in which $c_f =$ PHP 100 and $MO =$ PHP 995, and (3) a "high-incentive" shipping policy with a 50 percent higher shipping fee and minimum order requirement, i.e., c_f = PHP 150 and MO = PHP 1492.5. We simulate outcomes for the last year in our data, November 2015 to October 2016. In our simulation, we assume that consumers visit the website every other week. This is in line with the firm's browsing data, which show that 76% of visitors were last on the website within the previous two weeks. Moreover, we assume the shipping cost accruing to the seller to be PHP 100 per order.[8] This allows us to measure total profits, including both merchandise profits and profits from shipping.

Table 7 shows simulation results under the different shipping policies. We normalize all outcomes relative to the free shipping case, except for shipping revenue, which we benchmark against the baseline case. Compared with free shipping, the baseline contingent free shipping policy reduces the number of orders but increases the average order size conditional on purchase. The high-incentive shipping policy intensifies this tradeoff. Our simulation results further suggest that, for the focal retailer, the benefit from a larger average order size offsets the smaller number of orders, and raises merchandise revenue as well as merchandise profits. In addition, our simulation also shows that shipping fees provide substantial revenue for the retailer. As a result, higher shipping revenue under the high-incentive shipping policy further boosts retailer profits.

These changes in shipping policies also affect how consumers allocate their expenditure across categories. Simulated category expenditure shares in the lower half of Table 7 show that increasing the shipping fee and minimum order requirement results in more spending on dresses and shoes relative to other categories. As shown in Table 4, categories vary greatly in prices and unit margins; the high incentive shipping policy thus results in more spending on higher-priced, higher-margin categories, which provides a further boost to profits.

[8]While we do not possess detailed shipping cost data, our communication with the company confirms that this is a close approximate, and was in fact the main motivation for setting $c_f = 100$.

Table 7: Trade-off of the shipping policies

	Free shipping	**Baseline**	**High incentive**
The number of orders	1	0.936	0.905
Average order size	1	1.094	1.141
Merchandise revenue	1	1.024	1.033
Shipping revenue	–	1.000	1.892
Merchandise revenue + shipping revenue	1	1.037	1.057
Gross margin	1	1.000	1.001
Merchandise Profits	1	1.025	1.034
Total profits	1	1.124	1.212
Expenditure shares			
Accessories	1	0.992	0.988
Bags	1	1.000	0.998
Beauty	1	0.983	0.975
Bottoms	1	0.991	0.986
Dresses	1	1.005	1.006
Shoes	1	1.011	1.017
Tops	1	0.991	0.986
Underwear	1	0.973	0.960
Other	1	0.991	0.990

6.2 Shipping fees and product assortment

We now explore the interaction between product pricing and shipping fees. In particular, we examine whether a category price reduction has different impacts under different shipping policies. We simulate consumer demand for a 10% price cut for different categories under different shipping policies. For simplicity, we present the results for the two most purchased categories, shoes and tops, in Figure 5.

Under each shipping policy, a price cut for shoes, which is the most popular category, leads to more orders and also increases the average order size. Although both merchandise revenue and shipping revenue get a boost, the margin losses due to the price reduction decrease merchandise profits and total profits. We note general differences between tops and shoes: the same 10% price reduction for tops results in different patterns: (i) a smaller percentage increase in the number of orders, (ii) a decrease in the average order size, (iii) consequently, smaller impacts on merchandise

Figure 5: Effects of a 10 percent category price cut

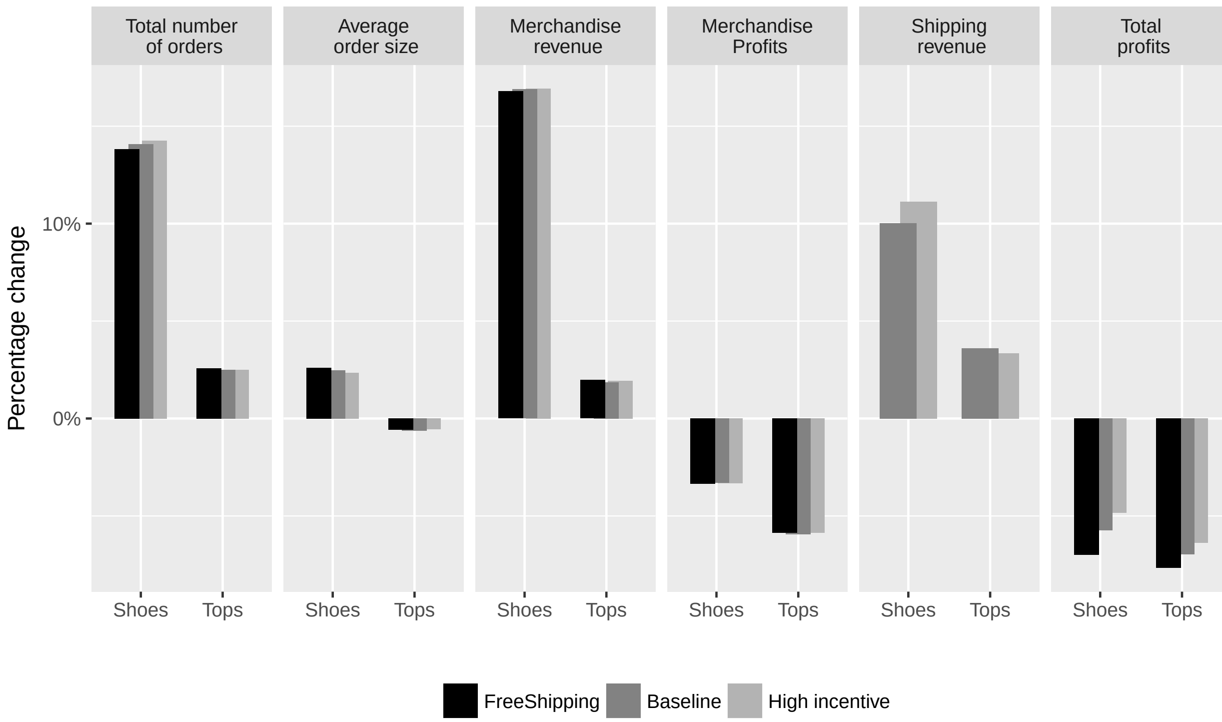

revenue and shipping revenue, and (iv) a larger decline in profit.

In both categories, the impact of changing prices varies by shipping policy. Under the baseline and high incentive shipping policies, reducing the average price of shoes is more effective at increasing conversion than in the free shipping condition, but has the opposite effect on the average order size. Shipping revenue is greater in the high incentive condition than in the baseline condition due to the increase in small orders. The effects of the same price reduction on merchandise profits and total profits decreases slightly in the contingent free shipping conditions relative to free shipping.

We observe that the differential price effects between different shipping policies depends on the category. If the retailer adopts a price reduction for the tops category, the price effects on the number of orders are roughly equal across different shipping policies, while the average order size is more negatively affected under the baseline shipping policy, and thus drags down total merchandise profits.

Inspecting Table 4, we see that shoes are not only more popular but also more expensive. In order

Figure 6: Effect of a 10 percent price cut (identical prices)

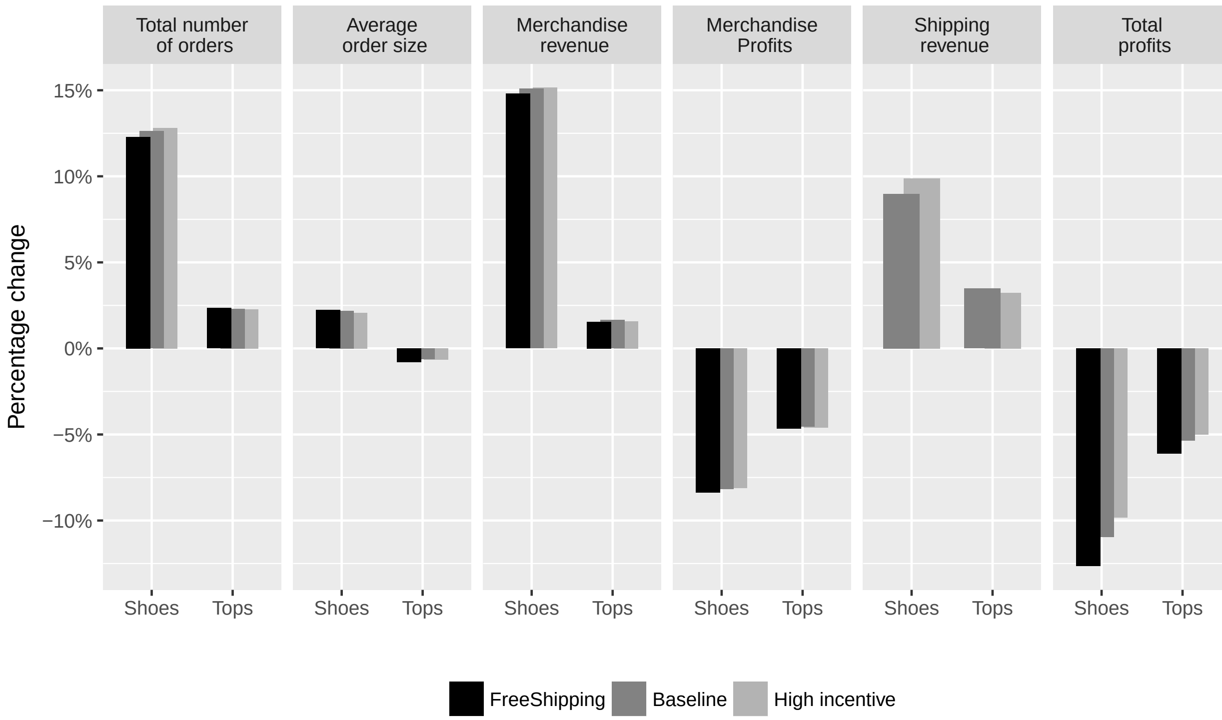

to identify the fundamental driver of differences in effects between categories, we conduct the same counterfactual simulation after equalizing the prices for shoes and tops. The percentage changes of various outcome variables are presented in Figure 6. The price effects shown in Figure 6 are very similar to those shown in Figure 5, suggesting that non-price attributes drive the observed differences in outcomes.

Our results show that shipping policies can moderate or amplify the impact of price changes on demand and profitability. Moreover, these interactions can vary between categories. Our simulations suggest that high-incentive shipping policies amplify the impact of price changes on conversion and order sizes for categories that are more attractive to consumers, net of prices.

6.3 Contingent free shipping and search costs

Our demand estimates show that it is costly for consumers to meet the minimum order threshold without substantially overshooting. We explore whether online retailers can benefit from lowering these search costs. In the following, we simulate consumer demand when search cost is decreased by 10%. Sellers can achieve this decrease by, for instance, providing product recommendations or making changes in product assortment.

Figure 7 presents the percentage changes in various outcome variables from a 10% decrease in search cost under the baseline and high-incentive shipping policies. In the baseline policy, lower search costs help consumers “top up” without significantly overshooting, and increases the number of orders, although the effect is small. The average order size drops slightly; however, the total number of consumers that qualify for free shipping increases. Hence, the retailer would collect lower shipping revenues.

Given a high-incentive shipping policy, we see that the same 10% decrease in search costs results in larger impacts on purchase behavior. In our first set of counterfactuals, we found that a high-incentive shipping policy leads to fewer orders but larger order sizes. Figure 7 suggests that lowering search costs can help mitigate this tradeoff, as it increases both the number of orders and the average order size. This results in an increase in merchandise revenue and merchandise profits. However, shipping revenue suffers as more consumers qualify for free shipping; this loss dominates gains from product sales.

Our findings point to a key instance in which easing search costs for consumers may hurt online sellers' profits. Contingent free shipping encourages consumers to increase their order sizes to qualify for free shipping; given the costly effort required to meet the minimum order requirement exactly, consumers may either forgo the incentive or substantially overshoot. When the seller lowers these local search costs, more consumers are able to qualify for free shipping without overshooting, and the net effect on profits is negative.

Figure 7: Effects of a 10 percent decrease of search cost.

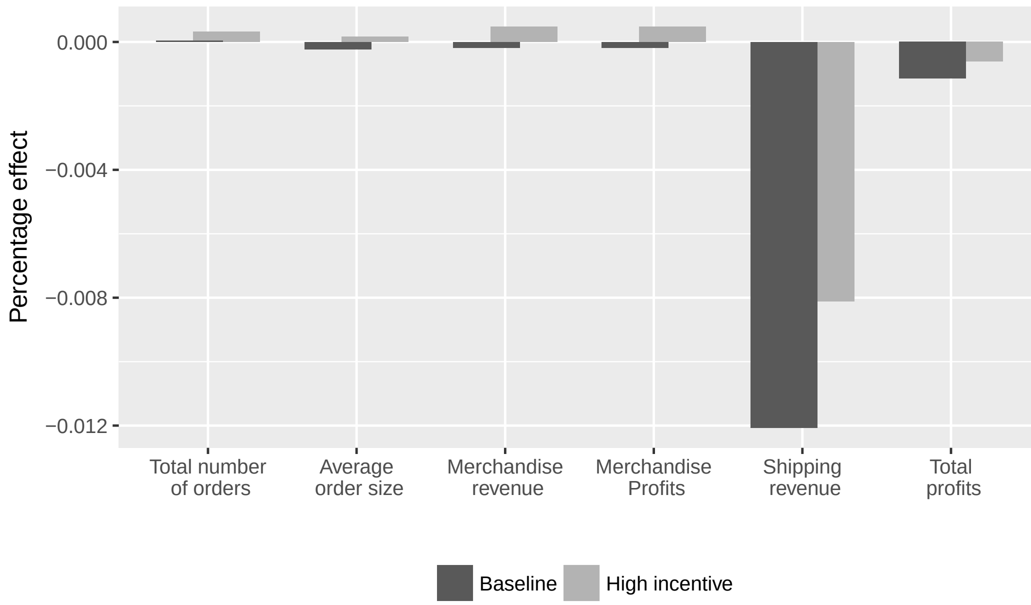

6.4 Free shipping subscription

Our last counterfactual is motivated by the observation that many major retailers, such as Amazon and Sephora, offer consumers memberships with a free shipping component. In these programs, consumers pay a fixed annual fee and enjoy free shipping for every order, regardless of basket size. Online retailers interested in launching similar shipping membership programs may seek to estimate how many consumers would sign up, how the program may influence purchase behavior, and whether such a program would be profitable.

We simulate outcomes given annual membership fee C_m. A consumer who purchases the membership has $I - C_m/12$ to spend on outside goods and makes purchase decisions on the retail website under the free shipping policy. In deciding whether to adopt the shipping membership, she compares the annual sum of the utility from adopting the membership against her utility given the baseline shipping policy.

Table 8 presents a set of outcome variables given membership fees. When the retailer charges PHP 200 for the membership, the adoption rate is 0.8 percent of potential customers. The higher the membership fee, the smaller the adoption rate. Importantly, consumers who adopt the membership are more likely to place larger-sized orders; hence, we see an increase in merchandise revenue and merchandise profits. Meanwhile, the retailer loses some shipping revenue with the program. Under our shipping cost assumption, the shipping membership program decreases total profits, but this profit gap narrows with higher membership fees.

Table 8: Shipping membership

	Baseline	**Membership fee**			
		200	300	400	500
Adoption rate	-	0.0081	0.0073	0.0070	0.0069
No. of orders	1	1.0074	1.0071	1.0073	1.0077
Average order size	1	1.0106	1.0118	1.0134	1.0146
Merchandising revenue	1	1.0180	1.0190	1.0208	1.0224
Merchandising profit	1	1.0180	1.0190	1.0208	1.0224
Shipping revenue	1	0.9717	0.9747	0.9753	0.9753
Membership revenue	-	1	1.3425	1.7126	2.1165
Total profits	1	0.8838	0.8856	0.8880	0.8903

Our findings point to an important tradeoff for online retailers who are interested in implementing such free shipping subscription programs. One the one hand, such a program increases the number of orders per customer as well as the average order size. On the other hand, these gains may not offset the loss in shipping revenue. In addition, under certain shipping cost conditions, additional orders may result in further losses.

7 Conclusion

Our findings illustrate the complex interplay between an online retailer's product assortment decisions and its shipping policies in determining purchase outcomes and profits. This interplay reflects the variety of variables affecting consumer utility. In choosing how to assemble their purchase baskets, consumers weigh not only the attractiveness of products but also how the addition of items may affect what they will owe in shipping fees. While motivated by a free shipping incentive to increase expenditure, consumers also face the costs of searching for products that will help meet the free shipping threshold with minimal overshooting.

We find that consumers are less sensitive to shipping fees than they are to product prices; nevertheless, contingent free shipping is a strong motivator for increasing average basket sizes. Furthermore, given differences between categories in marginal costs and utility parameters, the impact of price changes within a single category on conversion, basket sizes, and profits is highly dependent on the shipping policy. We find that it is costly for consumers to meet the minimum order threshold without substantially overshooting, and that it may not be profitable for online sellers to lower this search cost. We also show how a free-shipping membership can positively impact demand but have a negative impact on profitability.

Our research strategy introduces certain limitations. Because we adopt a static modeling approach, we may fail to capture intertemporal substitution resulting from free shipping incentives. We also lack clickstream data that may provide further insight on how shipping policies affect shopper search paths and cart abandonment decisions. Although we incorporate a rudimentary measure of shipping costs borne by the retailer, a more complete analysis of profitability would require access to detailed shipping, handling, and inventory cost data.

Given the centrality of shipping policies in the growing online retail sector, we believe that further research on this topic may provide valuable insights both on consumer behavior and firm policy. In particular, we feel that a fruitful area for future research lies in modeling multi-unit demand with discrete products, rather than the discrete-continuous approach we adopt here. Such an approach would provide the granularity needed to study how online recommendations and retargeting rules may be used together with free shipping thresholds, for example.

We also observe innovative pricing models with respect to shipping such as Amazon Prime Pantry, in which a box must be filled with pantry items before being eligible for delivery, or Jet.com, which offers progressive discounts for larger order sizes. These new pricing models provide both evidence of the importance of shipping policies for e-commerce managers and impetus for new research.

References

Becerril-Arreola, R., Leng, M., and Parlar, M. (2013). Online retailers' promotional pricing, free-shipping threshold, and inventory decisions: A simulation-based analysis. *European Journal of Operational Research*, 230(2):272–283.

Bertini, M. and Wathieu, L. (2008). Research note—Attention arousal through price partitioning. *Marketing Science*, 27(2):236–246.

Bhat, C. R. (2008). The multiple discrete-continuous extreme value (MDCEV) model: Role of utility function parameters, identification considerations, and model extensions. *Transportation Research Part B: Methodological*, 42(3):274–303.

Brown, J., Hossain, T., and Morgan, J. (2010). Shrouded attributes and information suppression: Evidence from the field. *The Quarterly Journal of Economics*, 125(2):859–876.

Brynjolfsson, E., Hu, Y., and Simester, D. (2011). Goodbye Pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management Science*, 57(8):1373–1386.

Chatterjee, P. (2011). Framing online promotions: Shipping price inflation and deal value perceptions. *Journal of Product & Brand Management*, 20(1):65–74.

Einav, L., Kuchler, T., Levin, J., and Sundaresan, N. (2015). Assessing sale strategies in online markets using matched listings. *American Economic Journal: Microeconomics*, 7(2):215–47.

Frischmann, T., Hinz, O., and Skiera, B. (2012). Retailers' use of shipping cost strategies: Free shipping or partitioned prices? *International Journal of Electronic Commerce*, 16(3):65–88.

Gümüş, M., Li, S., Oh, W., and Ray, S. (2013). Shipping fees or shipping free? A tale of two price partitioning strategies in online retailing. *Production and Operations Management*, 22(4):758–776.

Heiss, F. and Winschel, V. (2008). Likelihood approximation by numerical integration on sparse grids. *Journal of Econometrics*, 144(1):62–80.

Hendel, I. and Nevo, A. (2006). Measuring the implications of sales and consumer inventory behavior. *Econometrica*, 74(6):1637–1673.

Hossain, T. and Morgan, J. (2006). ... plus shipping and handling: Revenue (non) equivalence in field experiments on eBay. *Advances in Economic Analysis & Policy*, 5(2).

Jarvi, E. (2017). How to choose a shipping strategy for your online store. https://www.shopify.com/blog/8343330-how-to-choose-a-shipping-strategy-for-your-online-store. Accessed: 2017-10-23.

Jiang, Y., Shang, J., and Liu, Y. (2013). Optimizing shipping-fee schedules to maximize e-tailer profits. *International Journal of Production Economics*, 146(2):634–645.

King, G. and Zeng, L. (2001). Logistic regression in rare events data. *Political Analysis*, 9(2):137–163.

Leng, M. and Becerril-Arreola, R. (2010). Joint pricing and contingent free-shipping decisions in B2C transactions. *Production and Operations Management*, 19(4):390–405.

Lewis, M. (2006). The effect of shipping fees on customer acquisition, customer retention, and purchase quantities. *Journal of Retailing*, 82(1):13–23.

Lewis, M., Singh, V., and Fay, S. (2006). An empirical study of the impact of nonlinear shipping and handling fees on purchase incidence and expenditure decisions. *Marketing Science*, 25(1):51–64.

Morwitz, V. G., Greenleaf, E. A., and Johnson, E. J. (1998). Divide and prosper: consumers' reactions to partitioned prices. *Journal of Marketing Research*, pages 453–463.

Pinjari, A. R. and Bhat, C. (2010). A Multiple Discrete–Continuous Nested Extreme Value (MDC-

NEV) model: Formulation and application to non-worker activity time-use and timing behavior on weekdays. *Transportation Research Part B: Methodological*, 44(4):562–583.

Reagan, C. (2017). Think running retail stores is more expensive than selling online? Think again. https://www.cnbc.com/2017/04/19/think-running-retail-stores-is-more-expensive-than-selling-online-think-again.html. Accessed: 2017-10-23.

Shampanier, K., Mazar, N., and Ariely, D. (2007). Zero as a special price: The true value of free products. *Marketing Science*, 26(6):742–757.

Shriver, S. and Bollinger, B. (2015). Structural analysis of multi-channel demand.

Smith, A. S. (2018). NRF forecasts retail sales will increase between 3.8 and 4.4 percent in 2018. https://nrf.com/media/press-releases/nrf-forecasts-retail-sales-will-increase-between-38-and-44-percent-2018. Accessed: 2018-7-6.

Tyan, S. (2005). The Effects of Shipping Costs on Bidder Entry and Seller Revenues in eBay Auctions. *Senior Thesis, Department of Economics, Stanford University*.

Xu, J. (2016). *Empirical Studies In Online Retail Operations And Dynamic Pricing*. PhD thesis, University of Pennsylvania.

Yang, Y., Essegaier, S., and Bell, D. R. (2005). Free shipping and repeat buying on the internet: Theory and evidence. *Manuscript, Wharton School, University of Pennsylvania*.

Yao, Y. and Zhang, J. (2012). Pricing for shipping services of online retailers: Analytical and empirical approaches. *Decision Support Systems*, 53(2):368–380.

Appendix

Additional tables and figures

Figure A: Number of transactions per day

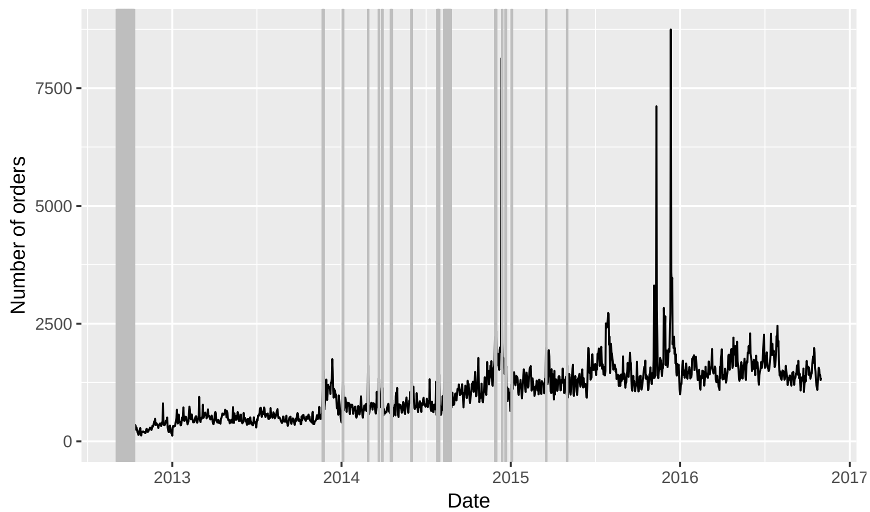

Note: Vertical gray lines indicate free shipping days.

Figure B: Number of items in order

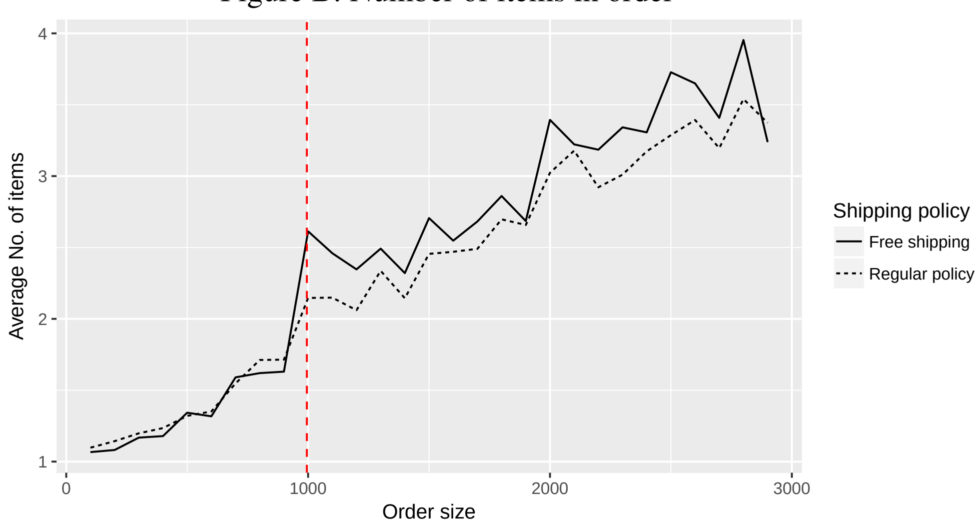

Figure C: Average prices within categories over time