

Hassle Costs and Price Discrimination: An Empirical Welfare Analysis[†]

By GUILLERMO MARSHALL*

This paper studies a market where soda is sold in both refillable and nonrefillable bottles. Purchasing refillables is inconvenient but cheaper. Using a discrete choice model, I find that price-sensitive customers put less weight on the inconveniences of purchasing refillables. This implies that a retailer can target lower prices to price-sensitive customers using the refillable segment. I evaluate the overall welfare consequences of this market segmentation and find that both customer welfare and profits would decrease (by 12.61 and 4.21 percent, respectively) if the refillables were removed, as there would be an important market-shrinkage effect. (JEL D22, L13, L25, L81)

Price discrimination is a widespread practice. Examples include student discounts at movie theaters, international-edition textbooks, and temporal sales at retail stores. The driving force of this practice is that customers are heterogeneous.

Practicing price discrimination requires firms to sort customers. Sorting customers can be straightforward in some cases (e.g., separating students from nonstudents using student ID cards) or challenging in others. To achieve sorting, firms often impose a hassle or effort cost to access a lower price. An example is when firms issue coupons. By varying both the conditions of use and the time burden involved in finding the coupon, firms can make sure that using the coupon is costly enough so that not every customer would be willing to use it.

A critical aspect of using hassle costs as a sorting device is understanding how the customers' price sensitivities and hassle-cost valuations are jointly distributed. In the previous example, the profitability of using a coupon relies on price-sensitive customers being the customers who find using the coupon attractive, as that would allow firms to expand coverage by charging a lower price to the more price-sensitive customers, without also having to lower the price for the least price-sensitive customers.

In this paper, I empirically study the welfare implications of market segmentation where hassle costs are used to sort consumers. While a monopolist is always better off under market segmentation, the overall welfare effect is theoretically ambiguous

*Department of Economics, University of Illinois at Urbana-Champaign, 214 David Kinley Hall, 1407 W. Gregory, Urbana, Illinois 61801 (e-mail: gmarshall@illinois.edu). I am grateful to Igal Hendel, Fernando Luco, Aviv Nevo, Álvaro Parra, Esteban Petruzzello, Tiago Pires, Rob Porter, Alberto Salvo, Anthony Wray, and three anonymous referees for comments and suggestions. The author declares that he has no relevant or material financial interests that relate to the research described in this paper.

[†]Go to <http://dx.doi.org/10.1257/app.20130046> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

and will depend on how market segmentation changes equilibrium quantity.¹ In support of the argument in the previous paragraph, a credible answer requires recovering customer preferences and, in particular, understanding how hassle-cost valuations and price sensitivities are jointly distributed.

The industry under study is the cola-drinks market in a developing country. In this particular market, customers can purchase a brand-size combination in a refillable bottle or a disposable bottle or both (henceforth, refillables and nonrefillables).² For a specific brand-size combination, purchasing the refillable format can be more than 20 percent cheaper. The lower prices, however, are not without cost. To purchase a refillable, a customer must either return a refillable bottle or invest in a new refillable bottle (i.e., pay a given amount in addition to the shelf price of the refillable product).

The availability of two different bottle formats causes customers to self-select. Given customer heterogeneity, retailers can adjust the price differences and the level of the hassle cost to maximize profits. As will be discussed in the next section, I find that the price differences between same brand-size combinations of different container formats reflect more than just cost differences, supporting the price discrimination hypothesis.

I use rich and unique data provided by a chain of supermarkets in a developing country. The data include customer demographics and customer-level information on all purchases that took place during a 91-week period. In addition, the data also include the wholesale prices of each product, which allows me to both implement tests for price discrimination as well as perform counterfactual pricing exercises without needing to impute wholesale price values.

My strategy in this study has two steps. The first step is to estimate preferences. To achieve this, I use a discrete choice framework to model the demand for soda. The demand model allows for customer heterogeneity in preferences, which depends both on observed and unobserved customer characteristics.

The second step is to simulate a scenario where refillables are removed from the market, thereby eliminating the ability of retailers to target lower prices to price-sensitive customers. This counterfactual scenario provides a benchmark for computing the welfare implications of price discrimination. Since the welfare analysis involves removing products from the market, I use a demand model that has the property that removing products from the choice set does not a priori affect welfare.³

From the estimates of the demand model, I find a negative correlation between the price coefficient and the hassle cost or taste for purchasing the refillable format. That is, price sensitive customers on average favor refillables more. This suggests that a retailer can exploit these customer differences and use the refillable segment to target lower prices to price-sensitive customers.

When simulating the removal of refillables, I find that both profits and customer welfare would decrease. I find that a retailer would see its profits drop by 4.21 percent

¹A monopolist is better off under price discrimination because the monopolist can always set the hassle cost equal to zero and set a uniform price. See Varian (1989) for a detailed discussion on the total welfare effect of price discrimination.

²Since I am considering a developing country, the preferences for recycling are a second-order concern.

³Models that include product-specific taste shocks with full support (e.g., logit model) do not have this property.

when losing its ability to both sell refillables and target lower prices to price-sensitive customers. Customers, on the other hand, would face an average welfare loss of almost 13 percent with the removal of refillables. Most of this welfare loss is a result of consumers leaving the market as a consequence of losing access to low-price soda.

Finally, to study how the welfare results rely on the negative correlation between price sensitivity and hassle costs, I manipulate the estimated distribution of preferences and repeat the welfare analysis. In support of the discussion above, I find that profits would suffer less with the removal of refillables if the coefficient of correlation was closer to zero (i.e., if the price-insensitive customers were not particularly deterred from purchasing refillables). This result suggests that the profit gains of serving customers with both formats depends on the nature of customer heterogeneity, as that dictates whether the seller can use the refillables to practice price discrimination.

The rest of the paper is organized as follows. The data and institutional background are discussed in Section II. The demand model is discussed in Section III. The estimation and the estimation results are discussed in Section IV. The counterfactual exercise that analyzes the welfare implications of market segmentation is discussed in Section V. Section VI concludes.

Literature Review.—This paper is related to several strands of the literature. The first is the literature studying the welfare implications of price discrimination. Empirical studies include Hendel and Nevo (2013), Leslie (2004), Villas-Boas (2009), and others. The closest work is Hendel and Nevo (2013). Hendel and Nevo (2013) use data on the soda industry in the United States to study how retailers use sales to intertemporally price discriminate among two types of consumers: storers and nonstorers. As in this paper, the authors argue that the retailer exploits multi-dimensional customer heterogeneity (i.e., in storage cost and price sensitivity) to sort customers based on price sensitivity.

Second, this paper is related to the literature studying the refillable bottle system in the beverage industries. Porter (1978) performs a social cost-benefit analysis of mandatory deposit on container deposits. Lesser and Madhavan (1987) perform a cost analysis of this same policy. While these articles consider environmental aspects, my analysis focuses on how a similar refillable-bottle system can enhance welfare when it enables price discrimination.

Finally, this paper is related to the literature on demand estimation. The model in this paper is similar to the one in Song (2012) and is based on ideas in McFadden (1981) and Berry and Pakes (2007). See Nevo (2011) and the references therein for an extensive literature review.

I. Data, Background, and Descriptive Analysis

A. Data

The data used in this paper were provided by a chain of supermarkets located in a large metropolitan area in a developing country. The sample period covers 91 consecutive weeks during years 2010 and 2011.

The data contain customer-level purchase and demographic information for about 50,000 randomly selected customers enrolled in the retail chain's loyalty program.⁴ The demographic data include the age and gender of each customer. The purchase data include records for all items that were purchased by each individual customer over the 91-week sample period. An observation in the purchase data is defined as a customer-trip-product combination and each observation includes a customer identifier, a trip identifier, the date of the trip, a product identifier, the transaction price, and the purchased quantity.

In addition, the data indicate the prices and wholesale prices of each of the products that was available in each store-week combination.⁵ These data allow for the reconstruction of the choice sets available to the customers on each trip to the store. Due to a privacy agreement with the supermarket chain, the location, the specific number of stores, and the wholesale prices will remain undisclosed.⁶

The sample I use in the analysis (both descriptive and structural) is restricted in two ways. First, I restrict the sample to the set of customers who purchased at most one bottle of cola during any given trip (71.16 percent of all customers). Restricting the sample in this way allows me to conduct the analysis while abstracting away from potential stockpiling incentives that arise in the context of storable-goods markets.⁷ Additionally, in order to abstract away from substitution between brands, I further restrict the sample to customers who always purchased the same brand of soda (e.g., always Diet Coke or always Coke Zero).⁸ This second restriction reduces the sample to 52.26 percent of all customers, representing approximately 700,000 shopping trips. Consequently, the results and analysis will be valid for nonstorer who were loyal to a brand of soda.

B. Background: Purchasing Refillables

Customers in this market can purchase a brand-size combination in a refillable bottle, or a disposable bottle, or both. To purchase in the refillable format, a customer must either return an empty refillable bottle of any brand-size or invest in a refillable bottle (i.e., pay a set amount, ranging from 60 cents to a dollar, in addition to the shelf price of the refillable product). As noted above, retailers do not require customers to return a refillable bottle that matches the brand-size of the product being purchased. That is, customers do not face compatibility issues once they have invested in a refillable bottle.

Customers return empty refillable bottles using reverse vending machines that deliver a proof of return. Vending machines are typically located in the parking lot or next to the supermarket entrance. In contrast to a deposit system, where a customer can recover the money that was paid for a bottle, a refillable bottle in this market is an illiquid investment.

⁴ About 80 percent of total sales are made by loyalty program members.

⁵ The wholesale price is a measure of the product replacement cost.

⁶ As a reference, the number of stores are in the order of several dozen.

⁷ See Hendel and Nevo (2006) for a detailed discussion.

⁸ Note that this restriction includes customers who purchased bottles of different size and format as long as they were of the same brand.

TABLE 1—SUMMARY STATISTICS

Brand	Size (L)	Refillable	Mean	SD	Min	Max	$\frac{\text{var}(p_{jt})}{\text{var}(p_{jst})}$	$\frac{\text{var}(p_{jst} - p_{jt})}{\text{var}(p_{jst})}$
<i>Panel A. Prices</i>								
Diet Coke	1.5	No	1.91	0.11	1.36	2.60	0.49	0.51
Diet Coke	1.5	Yes	1.49	0.11	1.32	1.95	0.52	0.48
Diet Coke	2.5	Yes	2.15	0.09	1.80	2.46	0.06	0.94
Diet Coke	3	No	3.11	0.18	2.06	3.68	0.33	0.67
Coke	1.5	No	1.86	0.11	1.34	2.60	0.41	0.59
Coke	1.75	No	2.08	0.22	1.01	2.17	0.73	0.27
Coke	2	Yes	1.79	0.10	1.34	2.58	0.27	0.73
Coke	2.5	Yes	2.06	0.08	1.78	2.46	0.13	0.87
Coke Zero	1.5	No	1.88	0.12	1.34	2.60	0.33	0.67
Coke Zero	1.75	No	2.15	0.07	1.91	2.17	0.34	0.66
Coke Zero	2	Yes	1.80	0.10	1.34	2.34	0.17	0.83
Coke Zero	2.5	No	2.64	0.16	2.06	3.00	0.39	0.61
Coke Zero	2.5	Yes	2.06	0.07	1.78	2.48	0.08	0.92
Coke Zero	3	No	2.91	0.14	2.06	3.66	0.29	0.71
<i>Panel B. Customer observables</i>								
Variable	Mean	SD	25th percentile	75th percentile	Min	Max		
Refillable purchases	0.29	0.42	0.00	0.67	0.00	1.00		
Age	37.64	15.01	26.00	47.00	18.00	88.00		
Mean trip expenditure	32.04	30.39	13.65	40.53	1.64	1,184.65		
Trip frequency	0.48	0.26	0.27	0.69	0.01	1.00		

Notes: Panel A was constructed using the weekly average price of each product over time and across stores. An observation is a product-store-week combination. Prices and trip expenditure are measured in US dollars. p_{jst} is the price of product j at store s in week t . p_{jt} is the average price of product j in week t . Panel B was constructed using customer observables. An observation is an individual customer. Mean trip expenditure is the customer's average expenditure across all trips to the store. Trip frequency is the ratio of weeks in which the customer visited a store over total weeks in the sample period. Refillable purchases is the ratio of purchases involving a refillable over total purchases involving soda.

C. Prices

Table 1 (panel A) shows descriptive statistics for the prices of each product. The table shows two patterns of price variation, which I discuss in the next paragraphs. The first one is that prices vary significantly within product. The second is that there are systematic price differences across formats even when conditioning on a brand-size combination.

I start studying the within-product price variation. Several sample statistics reported in Table 1 show the importance of this source of variation. For example, the maximum recorded price for a Diet Coke–1.5L–nonrefillable is almost twice its minimum price, or the coefficient of variation of a Coke–2L–refillable is almost 6 percent. However, interpreting these price differences is not straightforward, as they aggregate variation both over time and across stores. To understand the relative importance of these two contributing sources, I decompose the variance of prices using the following identity,

$$(1) \quad p_{jst} = p_{jt} + (p_{jst} - p_{jt}),$$

where p_{jst} is the price of product j at store s in week t , and p_{jt} is the average price of product j in week t . The first term on the right-hand side of (1) captures the time component of price, while the second term captures price differences across stores

in a given week. Using an argument analogous to the one given by Kaplan and Menzio (2014), one can show that the variance of p_{jst} is exactly equal to the sum of the variances of p_{jt} and $(p_{jst} - p_{jt})$. The results of this variance decomposition are displayed in the last two columns of Table 1 (panel A). The table shows that on average 32 percent of the price variation can be explained by the time component of prices (i.e., $\text{var}(p_{jt})$), while the remaining share of the price variation is explained by price differences across stores in a given week. These results show that both store and time variation are relevant for explaining the patterns in Table 1, but that most of the differences come from stores pricing the same products differently in a given week.

To complement this variance decomposition, Figure 1 displays the price evolution of the Diet Coke products in two stores over the sample period. Looking at how prices change over time can be helpful to understand whether the time variation of prices is given by the existence of sales or cost fluctuations or both. The figure shows that prices generally remain stable for several weeks, that sales are not frequent, and that price changes are mainly due to permanent price changes. This suggests that most of the price variation over time is given by cost fluctuations that affect the wholesale cost and that are passed on to price.⁹

I next study aspects of the price variation between formats. One can observe the average price difference between the nonrefillable and refillable formats (conditional on brand and size) by comparing the price difference of both formats of Diet Coke–1.5L and Coke Zero–2.5L in Table 1 (panel A). On average, the Diet Coke–1.5L–nonrefillable and the Coke Zero–2.5L–nonrefillable are, respectively, 42 and 58 cents more expensive than their same brand-size refillable alternatives.¹⁰ These differences imply that by purchasing Diet Coke–1.5L or Coke Zero–2.5L in the refillable format, a customer would save an average of almost 22 percent of the price of the nonrefillable alternative.

Figure 2 complements Table 1 as it displays the distribution of the price difference between both formats of Diet Coke–1.5L and Coke Zero–2.5L. The histogram shows that nonrefillables are always more expensive, which reflects that customers have to receive an incentive to bear the disamenities of purchasing refillables. The figure also shows that the price difference varies and can be as much as 90 cents.

Finally, I utilize the wholesale prices to study the extent to which the price differences in Figure 2 can be explained by cost differences. If in fact the price differences only reflect cost differences, one cannot claim that the retailer is practicing price discrimination. To study the nature of the price differences between formats, I use prices and wholesale prices to compute the price-cost margin and two markup measures for both formats of the Diet Coke–1.5L and Coke Zero–2.5L products (i.e., refillable and nonrefillable). I then run a separate difference in means test for both brand-size combinations, where the null hypothesis is that the difference in price-cost

⁹ Also noteworthy in Table 1 (panel A) and Figure 1 is that the price variation seems to be larger for nonrefillables. While the cost of the soda in a container is the same regardless of the type of container, the relative difference in the price variation between container formats may reflect that the production of a nonrefillable product is more plastic-intensive than that of a refillable product. Fluctuations in the price of plastic will affect the cost of producing a nonrefillable relatively more and, hence, will contribute to the difference in price variation across formats.

¹⁰ Using the prices of all products in Table 1 (panel A) and a hedonic price regression, I find that refillables are on average 55 cents cheaper than nonrefillables when controlling for brand, size, week, and store fixed effects.

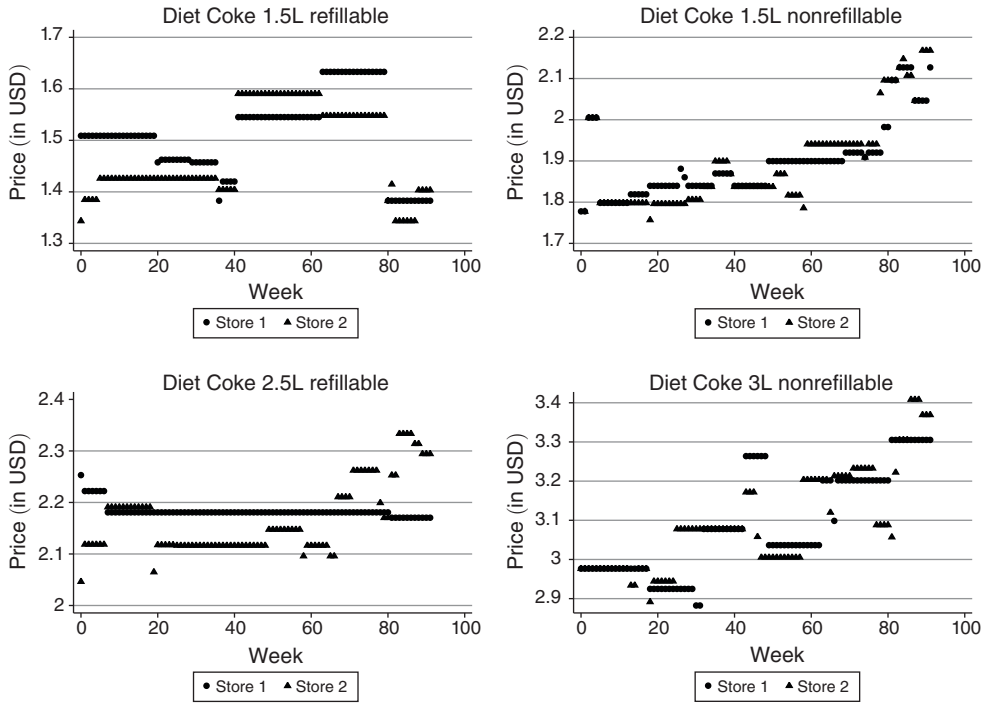


FIGURE 1. EVOLUTION OF PRICES BY PRODUCT

Notes: The figures were constructed using the weekly average price series of each product for two stores. See Table 1 for price summary statistics.

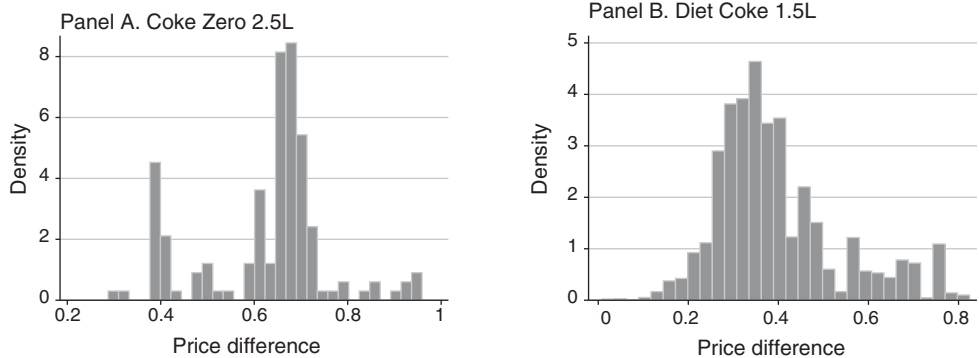


FIGURE 2. DISTRIBUTION OF PRICE DIFFERENCES ACROSS FORMATS CONDITIONAL ON BRAND SIZE

Notes: Price difference is defined as price nonrefillable—price refillable. The observations include store-week combinations in which both formats of the brand-size combinations were available.

margins (markups) can be explained by cost differences.¹¹ The results of these tests are presented in Table 2. The table shows that the null hypothesis—that price

¹¹ See Clerides (2004) for a discussion of these tests.

TABLE 2—ARE COST DIFFERENCES EXPLAINING THE PRICE DIFFERENCE BETWEEN FORMATS?
TESTING FOR MARGIN DIFFERENCES

Test	Coke Zero (2,500 cc) <i>t</i> -statistic	Diet Coke (1,500 cc) <i>t</i> -statistic
$H_0 : p_R - c_R = p_{NR} - c_{NR}$	38.2781	21.1465
$H_0 : \frac{p_R - c_R}{p_R} = \frac{p_{NR} - c_{NR}}{p_{NR}}$	21.2738	50.3785
$H_0 : \frac{p_R}{c_R} = \frac{p_{NR}}{c_{NR}}$	19.9471	49.9824

Notes: The tests performed are all two-sample *t*-tests, allowing for unequal variances. (p_R, c_R) and (p_{NR}, c_{NR}) are the price and wholesale prices of the refillable and nonrefillable versions of Diet Coke 1.5L or Coke Zero 2.5L, respectively, at a given store-week combination. The data used for these tests are the vector of prices and wholesale prices of each product across stores and weeks.

differences reflect only cost differences—is rejected in all tests, which provides evidence in favor of price discrimination.

D. Product Choice and Customer Sorting

Table 1 (panel B) shows observable characteristics of the customers in the sample. The average age of customers was 37.64 years, the average trip expenditure was 32.04 dollars, and the average trip frequency was once every two weeks. The average refillable-purchases ratio of customers was 0.10, which means that customers on average purchased a refillable product once every ten purchases of soda.

To understand how the observables in Table 1 (panel B) are related to the frequency with which customers purchased refillables (i.e., refillable purchases), I perform the following exercise. First, I divide the customers into four groups based on the refillable-purchases variable: (i) nonrefillable-loyal customers (i.e., refillable purchase intensity $\in [0, 0.01]$), (ii) low-intensity refillable customers (i.e., refillable purchase intensity $\in (0.01, 0.33]$), (iii) high-intensity refillable customers (i.e., refillable purchase intensity $\in (0.33, 0.99]$), and (iv) refillable-loyal customers (i.e., refillable purchase intensity $\in (0.99, 1]$).¹² Next, I run ordered probit regressions where the dependent variable is the refillable-intensity group number as defined above (i.e., $\{1, 2, 3, 4\}$).

The results of the ordered probit regressions are reported in Table 3. The table provides evidence of the existence of customer sorting between container formats. The regressions show that trip expenditure is negatively correlated with the intensity with which a customer purchases refillables. This finding suggests that customers that are less well-off on average favor refillables. The regressions also show that visiting the store more frequently is also positively correlated with purchasing refillables at a higher frequency.

Next, I study the importance of the refillable-nonrefillable price difference for format choice. For this, I consider individual choices in which a Diet Coke–1.5L of

¹²The threshold that divides customers between low- and high-intensity refillable customers is given by the median refillable purchase intensity (i.e. 0.33) among all customers with refillable purchase intensity $\in (0.01, 0.99)$. The results presented below are robust to varying this threshold.

TABLE 3—REFILLABLE PURCHASE INTENSITY ON CUSTOMER OBSERVABLES:
ORDERED PROBIT REGRESSIONS

	Refillable purchases category			
	(1)	(2)	(3)	(4)
Mean trip expenditure	-0.005*** (0.001)			-0.005*** (0.001)
Trip frequency		0.307*** (0.046)		0.244*** (0.047)
Age			-0.000 (0.001)	0.001 (0.001)
Observations	10,246	10,246	10,246	10,246
Pseudo R ²	0.006	0.002	0.000	0.007

Notes: Robust standard errors in parentheses. An observation is an individual customer. Refillable purchases category is defined to take one of four possible categories: (i) Refillable purchases ≤ 0.01 ; (ii) $0.01 < \text{Refillable purchases} \leq 0.33$; (iii) $0.33 < \text{Refillable purchases} \leq 0.99$; and (iv) Refillable purchases > 0.99 . Refillable purchases is the ratio of purchases involving a refillable over total purchases involving soda. See Table 1 for the definition of the remaining variables. Prices and trip expenditure are measured in US dollars.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

either refillable or nonrefillable format was purchased, and analyze how the choice of format is correlated with the price difference. In this exercise, I focus on customers who are either low-intensity or high-intensity refillable customers (see definition above). That is, I focus on the set of customers who alternate between formats and who are likely to be responding to the refillable-nonrefillable price difference.

Table 4 presents the results of this exercise. Specification 1 includes both low-intensity and high-intensity refillable customers and shows that the higher the price of the Diet Coke-1.5L-refillable relative to the Diet Coke-1.5L-nonrefillable, the lower the likelihood of a customer purchasing the refillable format. Specifications 2 and 3 repeat the exercise but restrict the sample to include only low-intensity and high-intensity refillable customers, respectively. As in the first specification, the price difference also affects the format choice in the expected way when restricting the sample of customers. What is interesting is that the point estimate of the price coefficient is relatively more negative for high-intensity refillable customers. This result suggests that customers who are on average more willing to bear the cost of purchasing in the refillable format are also more price sensitive. This relationship between hassle costs and price sensitivity is precisely the one that the retailer needs for the price discrimination technique to be profitable.¹³

II. Model

The model discussed in this section has the following features. First, the model allows for consumer heterogeneity in preferences. Second, the model incorporates an outside option of not purchasing soda, as understanding how market segmentation

¹³This is because price-insensitive customers would otherwise be exercising the option of a lower price.

TABLE 4—FORMAT CHOICE ON PRICE DIFFERENCES: OLS REGRESSIONS

Dependent variable:	Refillable choice		
	Low and high intensity (1)	Low intensity (2)	High intensity (3)
Sample of refillable customers:			
Price difference (Refillable – nonrefillable)	−1.503** (0.317)	−0.761** (0.378)	−2.043** (0.908)
Week FE	Yes	Yes	Yes
Store FE	Yes	Yes	Yes
Observations	546	380	166
R^2	0.384	0.418	0.797

Notes: Robust standard errors in parentheses. The table was constructed using individual choice decisions in which either a Diet Coke 1.5L of refillable or nonrefillable format was purchased when both were available. The variable refillable choice takes the value of one when the refillable format was chosen by the customer. Low (high) intensity refillable customers are customers with refillable purchases $\in (0, 0.33]$ ($\in (0.33, 0.99]$). Refillable purchases is the ratio of purchases involving a refillable over total purchases involving soda. Prices are measured in US dollars.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

affects quantity is only possible in a model that allows for customers exiting the market. Third, the model has the property that removing products from the choice set does not a priori affect welfare.

A. Demand

Consider an environment with T markets and N_t consumers in each market t , where a market is defined as a store-week combination.¹⁴ In market t , \mathcal{J}_t products are offered in addition to an outside option. A product is defined as a brand-size-format combination (e.g., Diet Coke–1.5L–refillable or Coke Zero–2.5L–nonrefillable).

I define $\mathcal{J}_{bt} \subset \mathcal{J}_t$ as the subset of products in market t that are of brand b . Given that the model is restricted to customers who are loyal to a brand, customer i in market t chooses among the set of products $\mathcal{J}_{it} \subset \mathcal{J}_t$ and an outside option, where $\mathcal{J}_{it} = \mathcal{J}_{b^*(i)t}$, and where $b^*(i)$ is customer i 's favorite brand.

I make the following assumptions regarding the information structure and the timing of the decisions:

ASSUMPTION 1:

- (i) Every customer owns a refillable bottle.
- (ii) Customers are informed about the choice set, prices of each product, and their taste shocks before visiting the store.

¹⁴The assumption that each market is served by a unique retailer can be justified by the fact that prices are found to be very similar for nearby stores of the same chain or by arguing that prices in a single product category are not pivotal for the decision of which store to visit.

These assumptions, while strong, simplify the analysis as they imply that a customer will never purchase a refillable without having brought a refillable bottle to the store. The reason is that, since the relevant information for the decision of which product to purchase is available before visiting the store, the consumer carries the refillable bottle to the store if and only if the customer decided in advance to purchase a refillable.¹⁵

I define consumer i 's indirect utility for product $j \in \mathcal{J}_{it}$ as

$$u_{ijt} = \alpha_i(I_i - p_{jt}) - \gamma_i r_j + \mathbf{x}'_j \beta + \varepsilon_{ib^*(i)t},$$

where p_{jt} is the price of product j in market t , I_i is consumer i 's income, r_j is a dummy that takes the value of one if the product is a refillable, and \mathbf{x}_j are observed and unobserved time-invariant characteristics of product j . The observed product characteristics include both container size and brand dummies. $\varepsilon_{ib^*(i)t}$ is a taste shock specific to customer i 's favorite brand $b^*(i)$.

Note that since the taste shock, $\varepsilon_{ib^*(i)t}$, is brand-specific rather than product-specific, the taste shock does not affect the customers' comparison of same-brand products. Consequently, a (α_i, γ_i) -type customer will never purchase a product that is dominated by another same-brand product. As will be discussed below, this implies that the model has the property that the removal of a product does not a priori affect welfare.¹⁶ Other models that include product-specific taste shocks with full support (e.g., logit demand model) do not share this property, as taste shocks always make purchasing an arbitrary good a positive-probability event.

In a similar way, I define the indirect utility that consumer i obtains from the outside option as

$$u_{i0t} = \alpha_i I_i + \varepsilon_{i0t}.$$

The model incorporates customer heterogeneity through two sources. The first is through the two taste shocks that each customer faces: one specific to the customer's favorite brand and one specific to the outside option. The second is through differences in the taste for two product attributes: the price and the disamenities attached to the refillable format. I define consumer i 's marginal utility of income and taste for the disutility attached to the refillable format as

$$(2) \quad \alpha_i = \alpha_1 + \mathbf{v}'_i \alpha_2 + \eta_{it}, \text{ and } \gamma_i = \gamma_1 + \mathbf{v}'_i \gamma_2 + \mu_{it},$$

respectively, where \mathbf{v}_i are observed customer characteristics, while η_{it} and μ_{it} are unobserved customer characteristics. I assume that η and μ are independent across customers and time and are distributed according to the normal distributions,

¹⁵ A more elaborate model would involve a two-stage decision: (i) decide whether to take the bottle to the store; and (ii) conditional on the first-stage decision, and after observing prices and product availability at the store, make the product choice.

¹⁶ In a discrete choice model, a customer's welfare decreases when a product that is purchased by the customer with positive probability is removed from the choice set.

$N(0, \sigma_\eta^2)$ and $N(0, \sigma_\mu^2)$, respectively. I also assume that η and μ are both independent of the idiosyncratic taste shocks, ε .

Given the specification of the model, the probability that customer $i \in N_t$ purchases good j is given by

$$s_{ijt}(\mathbf{p}_t, \mathbf{x}, \mathbf{v}_i; \theta) = \int 1\{u_{ijt} \geq u_{ikt}, \forall k \mid \mathbf{p}_t, \mathbf{x}, \mathbf{v}_i, \varepsilon, \mu, \eta; \theta\} dH(\varepsilon)dG(\mu)dF(\eta),$$

where θ is the vector of parameters that enter the indirect utility functions.

I assume that the idiosyncratic shocks $\{\varepsilon_{ib^*(i)t}, \varepsilon_{i0t}\}$ are independent across customers, brands, and time, and are distributed according to a Type 1 extreme value distribution. This implies that the probability above can be written as

$$(3) \quad s_{ijt} = \int \frac{\exp\{\max_{k \in \mathcal{J}_{b^*(i)t}} \{-\alpha_i p_{kt} - \gamma_i r_k + \mathbf{x}'_k \beta\}\}}{1 + \exp\{\max_{k \in \mathcal{J}_{b^*(i)t}} \{-\alpha_i p_{kt} - \gamma_i r_k + \mathbf{x}'_k \beta\}\}} \times 1\{u_{ijt} \geq u_{ikt}, \forall k \in \mathcal{J}_{b^*(i)t}\} dG(\mu)dF(\eta),$$

for $j \in \mathcal{J}_{it}$. The first term in expression (3) captures the probability that customer i with unobserved characteristics (μ_i, η_i) chooses among the products of his favorite brand $b^*(i)$, while the second term restricts the domain of integration to the region in which the customer chooses product j conditional on the event that brand $b^*(i)$ is chosen.

Note that given these assumptions, conditional on choosing brand $b^*(i)$, the only product chosen with positive probability by customer i , with unobserved characteristics (μ_i, η_i) , is the one that gives the customer the highest indirect utility among all options of brand $b^*(i)$. Adding products of brand $b^*(i)$ that are worse in terms of indirect utility would not change customer i 's decision.

B. Customer Welfare

For the analysis below, I use the inclusive value (or expected maximized utility) as the measure of customer welfare. This measure is defined as

$$(4) \quad E \left[\max_{k \in \mathcal{J}_{b^*(i)t} \cup \{0\}} u_{ijt} \right] \\ = \int \int \ln \left(1 + \exp \left\{ \max_{h \in \mathcal{J}_{b^*(i)t}} \{-\alpha_i p_{ht} - \gamma_i r_h + \mathbf{x}'_h \beta\} \right\} \right) dG(\mu)dF(\eta) + \gamma,$$

where γ is Euler's constant.

As can be noted from (4), a customer's welfare is dictated by the customer's preferred option in choice set \mathcal{J}_{it} . Since there are no product-specific taste shocks,

a customer of type (μ, η) will always choose the same product, as the customer's preferred product is a degenerate random variable.¹⁷ As a consequence, changing the choice set \mathcal{J}_{it} of a customer of type (μ, η) will affect the customer's welfare if and only if the change in the choice set implies a change in the customer's preferred option. As discussed above, this implies that a customer's welfare is not a priori affected by changes in the customer's choice set.

Using (4), I define a monetary measure of how customer i 's welfare is affected by a change from $(\mathbf{p}_t, \mathcal{J}_{it})$ to $(\mathbf{p}'_t, \mathcal{J}'_{it})$,

$$(5) \quad \Delta\omega_{it}(\mathbf{p}_t, \mathcal{J}_{it}, \mathbf{p}'_t, \mathcal{J}'_{it}) \\ = \int \frac{1}{\alpha_i} \left[\ln \left(1 + \exp \left\{ \max_{j \in \mathcal{J}_{b^*(i)t}} \{ -\alpha_i p_{jt} - \gamma_i r_j + \mathbf{x}'_j \beta \} \right\} \right) \right. \\ \left. - \ln \left(1 + \exp \left\{ \max_{j \in \mathcal{J}'_{b^*(i)t}} \{ -\alpha_i p'_{jt} - \gamma_i r_j + \mathbf{x}'_j \beta \} \right\} \right) \right] dG(\mu) dF(\eta).$$

C. Identification

For a given vector of prices and price coefficient α , the model interprets heterogeneity in the propensity to purchase refillables as heterogeneity in the disutility attached to purchasing refillables, γ .

Both brand dummies and the price coefficient α conflict in their effect on the overall market share of the inside options, but only the latter is able to capture how the probability of purchasing an inside option varies with the level of prices. How a customer changes his probability of purchasing an inside option as the vector of prices varies over time is what identifies α . The brand dummies, instead, are identified by the average propensity of purchasing an inside option relative to the outside option.

In this way, both cross-section and within-customer variation in choices play critical roles for identification. The cross-section variation in choices helps identify the disutility attached to refillables, while the latter helps identify how these choices are affected by price levels.

A concern with the above arguments is that prices may be responding to unobserved factors that may also affect the demand for a given product or brand. Advertising is an example of such a factor. Not controlling for these unobserved factors may result in inconsistent estimates of customer price sensitivity. To address these concerns, I use a control function approach in the estimation (Petrin and Train 2010).

¹⁷In contrast, with product-specific taste shocks, the customer's preferred product would be a nondegenerate option (e.g., the logit model).

III. Estimation

The estimation procedure I use for estimating the demand model is simulated maximum likelihood. Recall from Section II that the indirect utility that consumer i derives from purchasing product j is given by

$$u_{ijt} = \alpha_i(I_i - p_{jt}) - \gamma_i r_j + \mathbf{x}'_j \beta + \varepsilon_{ib^*(i)t},$$

where

$$\alpha_i = \alpha_1 + \mathbf{v}'_i \alpha_2 + \eta, \text{ and } \gamma_i = \gamma_1 + \mathbf{v}'_i \gamma_2 + \mu,$$

with η and μ being independent draws from $N(0, \sigma_\eta^2)$ and $N(0, \sigma_\mu^2)$, respectively. As discussed above, $(p_{jt}, r_j, \mathbf{x}_j, \mathbf{v}_i)$ are data, while $\theta = (\alpha_1, \alpha_2, \sigma_\eta, \gamma_1, \gamma_2, \sigma_\mu, \beta)$ is a vector of parameters to be estimated.

For the purposes of the estimation, I simulate the choice probabilities defined in (3) using Monte Carlo integration,

$$s_{ijt}(\mathbf{p}_t, \mathbf{x}, \mathbf{v}_i; \theta) = \frac{1}{D^2} \sum_{s=1}^D \sum_{r=1}^D \int 1\{u_{ijt} \geq u_{ikt}, \forall k \in \mathcal{J}_{b^*(i)t} \cup \{0\} | \mathbf{p}_t, \mathbf{x}, \mathbf{v}_i, \varepsilon, \mu_s, \eta_r; \theta\} dH(\varepsilon),$$

where $\{\mu_s\}_{s=1}^D$ and $\{\eta_r\}_{r=1}^D$ are independent draws from $N(0, \sigma_\mu^2)$ and $N(0, \sigma_\eta^2)$, respectively. Using these probabilities, I define the simulated likelihood function as

$$(6) \quad \mathcal{S}\mathcal{L}_N(\theta) = \sum_{t \in T} \sum_{i \in N_t} \sum_{j \in \mathcal{J}_i \cup \{0\}} d_{ij} \log(s_{ijt}),$$

where the binary variable d_{ij} is equal to one if consumer i in market t chose option j . As discussed above, the choice of the outside option is observed in the data and occurs when a trip is completed without the purchase of a cola. By definition,

$$\sum_{j \in \mathcal{J}_i \cup \{0\}} d_{ij} = 1.$$

The simulated maximum likelihood estimator is given by

$$\hat{\theta} = \arg \max_{\theta} \mathcal{S}\mathcal{L}_N(\theta).$$

A. Control Function

As discussed in the previous section, there is a concern that prices are correlated with unobserved factors, ϕ , that may also affect product choice. To address the issue of such unobserved factors, I also estimate the model using a control function

TABLE 5—SIMULATED MAXIMUM LIKELIHOOD ESTIMATES

	No control function		Control function	
	Estimate	SE	Estimate	SE
Price	-6.1548	0.0417	-8.1531	0.0076
Price × trip frequency	-1.6384	0.0716	-2.0646	0.0317
Price × mean expenditure	0.0132	0.0010	0.0150	0.0001
Price × age	0.0063	0.0002	0.0107	0.0001
Refillable	-1.6305	0.0300	-1.0038	0.0084
Refillable × trip frequency	0.6038	0.0234	0.7092	0.0061
Refillable × mean expenditure	-0.0066	0.0003	-0.0104	0.0001
Refillable × age	0.0055	0.0002	-0.0037	0.0001
Container size (L)	1.2878	0.0184	0.7063	0.0070
Diet Coke	-0.0418	0.0017	-0.0407	0.0003
Coke	0.2076	0.0112	0.1765	0.0015
Coke Zero	-0.1622	0.0056	-0.1940	0.0040
CF parameter	—	—	-0.0033	0.0001
σ'_{η}	1.1966	0.0050	1.6015	0.0015
σ'_{μ}	0.7208	0.0154	1.0071	0.0050
$-\frac{1}{N} \mathcal{S}\mathcal{L}_N(\hat{\theta})$	0.2602		0.2506	
Number of choices (N)	671,086		671,086	

Notes: Standard errors were computed using the bootstrap. Mean trip expenditure is the customer's average expenditure across all trips to the store. Prices and trip expenditure are measured in US dollars. Trip frequency is the ratio of weeks in which the customer visited a store over total weeks in the sample period (see Table 1). The parameters σ'_{η} and σ'_{μ} define the standard deviations of the distributions of the random coefficients on price and refillable disamenities: $\sigma_{\eta} = \exp\{\sigma'_{\eta}\}$ and $\sigma_{\mu} = \exp\{\sigma'_{\mu}\}$. CF parameter is the control function parameter.

approach (Pettrin and Train 2010). The identifying assumption behind this approach is that prices and taste shocks $\varepsilon_{ib(j)t}$ are independent once one conditions on the unobserved factors, ϕ . To recover a measure of these unobserved factors, I use the following specification for the pricing function:

$$p_{jt} = \delta_1 c_{jt} + \mathbf{x}'_j \delta_2 + \phi_{jt},$$

where c_{jt} is the wholesale cost of product j in market t , \mathbf{x}_j is a vector of product characteristics, and ϕ_{jt} are the unobserved factors that affect price and may also affect the choice of product j in market t .

Given the estimates of ϕ_{jt} , I specify the control function as a linear function of ϕ_{jt} , $\text{CF}(\phi_{jt}; \lambda) = \lambda \phi_{jt}$. The indirect utility function then becomes,

$$u_{ijt}^{\text{CF}} = \alpha_i(I_i - p_{jt}) - \gamma_i r_j + \mathbf{x}'_j \beta + \lambda \phi_{jt} + \varepsilon_{ib^*(i)t}$$

Under the identifying assumption that p_{jt} and $\varepsilon_{ib^*(i)t}$ are independent conditional on ϕ_{jt} , the simulated maximum likelihood estimator of θ is consistent.

B. Estimation Results

Table 5 reports the estimates of the model. The first and second columns report the estimates with and without the control function, respectively.

The results show a negative effect of price on the demand for soda, which varies with observed and unobserved customer characteristics. Customers who on average

spend more during each trip and who are older, are found to be less price sensitive. Customers who visit stores more frequently are on average more price sensitive. As can be noted, the sign of the price coefficients and the relative magnitude of these coefficients are consistent across specifications, though the coefficients are larger in absolute value in the control function specification.

The results also show that customers on average derive negative utility from the disamenities attached to the refillable format. Again, there is customer heterogeneity in how customers value refillables. Both specifications show that customers who on average spend more and visit the store less often dislike refillables more (all else held equal). Note that these relationships are consistent with the results in Table 3, where I find that the extent to which a customer purchases refillables is positively correlated with trip frequency and negatively correlated with average trip expenditure.

Figure 3 shows the joint distribution of the price coefficient and refillable coefficient of all customers (excluding unobserved heterogeneity). As can be seen in the figure, both the price and refillable coefficients are estimated to be negative for all customers, which implies that refillables must be priced lower in order for customers to purchase this format.

Figure 3 also shows a negative correlation between the price and refillable coefficients, meaning that customers who are more price sensitive are also customers who value less negatively the refillable format. This relationship is in line with the results in Table 4, where I find evidence consistent with a higher price sensitivity in customers who choose the refillable format more frequently. As will be discussed in the next section, this relationship implies that conditional on making a purchase, price sensitive customers are on average more likely to purchase refillables.

Finally, Table 5 also shows that customers on average value favorably larger containers and that customers who purchase Coke on average purchase soda more frequently than customers who purchase Diet Coke or Coke Zero.

IV. Market Segmentation Analysis

In this section, I evaluate the welfare consequences of forcing sellers to serve customers with the nonrefillable segment only. I perform this analysis in three steps. In the first step, I compute the optimal prices under the status quo (i.e., both formats available) using the estimates of the model (with control function).¹⁸ Using these prices and the preferences of customers, I compute profits and customer welfare.

In the second step, I assume that all refillable products are removed from the market. Given the restricted choice set, I compute the new optimal prices and use these to compute profits and customer welfare. I then proceed to compare how profits and customer welfare change when moving from the status quo to the case when refillables are removed from the market. This exercise allows us understand how welfare changes when the seller loses the ability to practice price discrimination.

In the final step, I study how the decrease in profits that results from the removal of refillables varies depending on how well the retailer can sort customers using the

¹⁸The attributes of the products are as displayed in Table 1, and the wholesale prices of each product are assumed to take the median values across the sample.

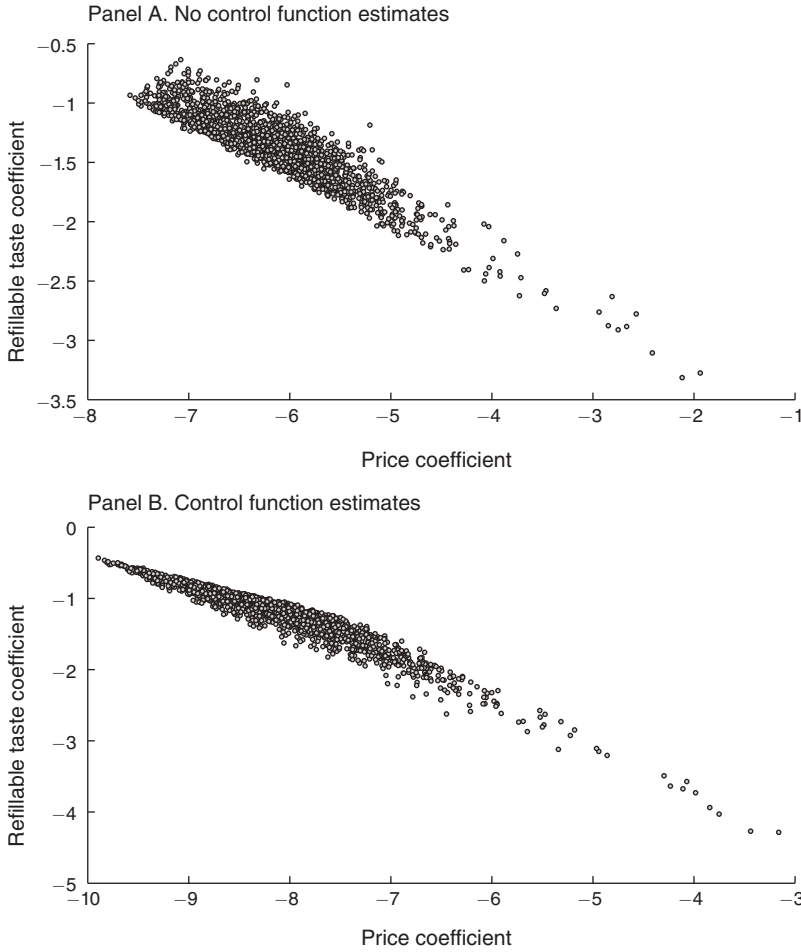


FIGURE 3. JOINT DISTRIBUTION OF COEFFICIENTS: PRICE AND REFILLABLE TASTE COEFFICIENTS

Notes: The coefficients that are displayed are the mean values of the consumers’ distributions of random coefficients. The price coefficient is $-\alpha_i$, that is, the negative of the marginal utility of income.

refillable segment. This analysis helps explain the role of customer sorting for the profitability of serving customers with both refillables and nonrefillables.

A. Status Quo

I start analyzing the optimal prices and market shares under the status quo (i.e., both formats available). Optimal prices are the solution to the problem,

$$(7) \quad \max_{\{p_j\}_{j \in \mathcal{J}_b}} \sum_{i \in N} \sum_{j \in \mathcal{J}_b} (p_j - c_j) s_{ij}(\mathbf{p}),$$

where p_j and c_j are the price and wholesale price of product j , respectively, and where s_{ij} is the probability that customer i purchases product j (see the definition in (3)).

TABLE 6—PRICES, MARKET SHARES, AND WELFARE MEASURES BEFORE AND AFTER REMOVAL OF REFILLABLES

Size (L)	Refillable	Before removal		After removal		Welfare change	
		Share	Price	Share	Price	CS	Profits
<i>Panel A. Diet Coke</i>							
1.5	No	0.0120	2.11	0.0560	2.17	—	—
1.5	Yes	0.0533	1.62	—	—	—	—
2.5	Yes	0.0008	2.13	—	—	—	—
3	No	0.0049	2.89	0.0052	2.88	—	—
Total:		0.0709 ^a	—	0.0612 ^a	—	-5.23%	-9.72%***
<i>Panel B. Coke</i>							
1.5	No	0.0001	2.17	0.0021	2.32	—	—
1.75	No	0.0394	2.13	0.0583	2.13	—	—
2	Yes	0.0252	2.05	—	—	—	—
2.5	Yes	0.0039	2.15	—	—	—	—
Total:		0.0686 ^a	—	0.0604 ^a	—	-14.07%*	-2.46%**
<i>Panel C. Coke Zero</i>							
1.5	No	0.0273	2.17	0.0541	2.17	—	—
1.75	No	0.0000	2.36	0.0005	2.32	—	—
2	Yes	0.0324	2.01	—	—	—	—
2.5	No	0.0009	2.82	0.0007	3.03	—	—
2.5	Yes	0.0023	2.13	—	—	—	—
3	No	0.0034	5.12	0.0032	5.38	—	—
Total:		0.0663 ^a	—	0.0585 ^a	—	-16.88%***	-3.59%
<i>Panel D. Overall</i>						-12.61%**	-4.21%**

^a Denotes that the difference between the labeled values is statistically significant at the 5 percent level.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Table 6 shows the prices and market shares before refillables are removed from the market. The table shows that the prices for the Diet Coke–1.5L–refillable and Coke Zero–2.5L–refillable products are about 50 and 60 cents lower than the prices for the Diet Coke–1.5L–nonrefillable and Coke Zero–2.5L–refillable products, respectively, which reflects both that sellers have to give customers incentives to choose refillables over nonrefillables and that sellers use refillables to target lower prices to price-sensitive customers. The table also shows that smaller containers (i.e., 1.5L and 1.75L) are most popular and that with one exception the optimal prices predicted by the model are in the range of the prices that are observed in the data (see Table 1).

To complement Table 6, Figure 4 shows the joint distribution of a customer's probability of choosing a Diet Coke–refillable good (conditional on purchase) and (mean) price coefficient. This figure shows how the relationship between a customer's price sensitivity and taste for the disamenities attached to the refillable format (see Figure 3) allows the retailer to sort customers by price sensitivity when both formats are on the market.¹⁹

¹⁹ Similar figures for Coke and Coke Zero show qualitatively identical results and are available upon request.

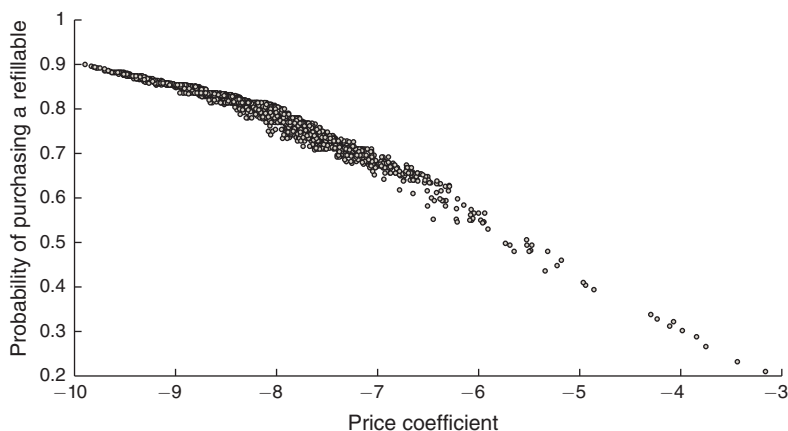


FIGURE 4. CUSTOMER SORTING:
PROBABILITY OF PURCHASING DIET COKE REFILLABLES (*conditional on purchase*) ON PRICE SENSITIVITY

Notes: The probability is the model's prediction of the probability of purchasing a refillable conditional on purchasing an inside option. The price coefficients in the figure are customer averages over the unobserved customer characteristics.

B. Retailer Profits

I next analyze how profits and welfare change when refillables are removed from the market. After the removal of refillables, the retailer sets prices by solving the problem in (7), but with a set of products that excludes refillables. Prices, market shares, and profits after the removal of refillables are displayed in Table 6.

As refillables are removed from the market, the composition of customers purchasing nonrefillables changes. The change in customer composition has two effects on how the retailer prices nonrefillables. On the one hand, as the more price-sensitive customers relied on the low-price refillables to purchase soda, the retailer has incentives to lower the price of the nonrefillables to discourage the price-sensitive customers from exiting the market. On the other hand, customer sorting in the model is not perfect (see Figure 4), meaning that some price-insensitive customers were purchasing refillables when refillables were available. This implies that with the removal of refillables, an increased number of price-insensitive customers would purchase nonrefillables, which gives the retailer an incentive to increase the price of nonrefillables.

The prices that the retailer sets after the removal of refillables capture both these effects. For instance, the prices of Diet Coke-3L-nonrefillable and Coke Zero-1.75L-nonrefillable decrease by 1 and 4 cents, respectively, which reflects that the effect of making these products more accessible to price-sensitive customers dominates the effect of serving proportionally more price-insensitive customers. On the contrary, the prices of Diet Coke-1.5L-nonrefillable and Coke-1.5L-nonrefillable increase by 6 and 15 cents, respectively, which reflects instead that the effect of serving proportionally more price insensitive customers dominates.

The market shares are affected both by the removal of refillables and the price changes. With the removal of refillables, customers are left with the relatively more expensive options, which affects the likelihood of a customer purchasing soda. As the

table shows, the total market shares drop for all three brands. The decrease in market share captures that the price decreases in some of the nonrefillables are unable to compensate for the removal of the affordable refillable options, which forces some customers to exit the market. The highest decrease in market share is in the Diet Coke products, where the overall market share drops from 7.09 percent to 6.12 percent with the removal of refillables. As will be discussed in the next subsection, price-sensitive customers contribute proportionally more to this drop in the market share.

Finally, with the removal of refillables, profits are found to drop by more than 4 percent. The profit decrease is highest for the Diet Coke products (almost 10 percent), which is the brand that suffers the highest decrease in market share. The drop in profits captures that the retailer loses the ability to target lower prices to price sensitive customers, reflected in the drop in quantity.²⁰

C. Customer Welfare

I next study how customer welfare changes once refillables are removed from the market. After the removal of refillables, the welfare of customers is affected as customers are both left with the relatively more expensive options and faced with price changes. To compute how these factors affect customer welfare, I compute the change in customer welfare before and after the removal of refillables by using the expression in (5). As Table 6 shows, removing refillables implies that customer welfare decreases by almost 13 percent. This effect on customer welfare parallels the effect on market share discussed above.

Figure 5 shows the joint distribution of the welfare change with the removal of refillables (in percent) and the (mean) price coefficient for Diet Coke customers. The figure shows that most customers are worse off with the removal of refillables, with an average welfare change of -5.24 percent. The figure shows that the effect is proportionally higher for more price-sensitive customers, which captures that quantity decreases the most for price sensitive customers as a consequence of cheaper refillable options no longer being available.²¹

The effect on profits and customer welfare imply that removing refillables harms overall welfare or, put differently, that market segmentation is welfare enhancing.

D. Profits and Customer Sorting

Finally, I analyze how the decrease in profits that results from the removal of refillables varies depending on how well the retailer can sort customers using the refillable segment. This analysis will help clarify the role of customer sorting for the profit gains of serving customers with both bottle formats. To perform this analysis, I manipulate the estimated distribution of price and refillable coefficients and repeat the welfare analysis above for a series of counterfactual distributions of preferences.

²⁰The removal of refillables may imply cost savings associated with the handling of the refillable bottles. In this sense, the profit gains of price discrimination that I report may be overestimated.

²¹Similar figures for Coke and Coke Zero show qualitatively identical results and are available upon request.




FIGURE 5. THE EFFECT OF REMOVING REFILLABLES ON WELFARE: DIET COKE PRODUCTS

Notes: Welfare change (in percent) is measured based on expression (4). The price coefficients in the figure are customer averages over the unobserved customer characteristic.

The counterfactual distributions of price and refillable coefficients are given by pairs $(\alpha_i, \gamma_i^{\text{Counterfactual}}(\rho))$, where α_i takes the value that was estimated from the model for customer i and $\gamma_i^{\text{Counterfactual}}(\rho)$ takes a value that is given by

$$(8) \quad \gamma_i^{\text{Counterfactual}}(\rho) = \underline{\gamma} + \rho(\gamma_i - \underline{\gamma}),$$

where $\underline{\gamma} = \min_{i \in I} \gamma_i$ (i.e., the lowest value among the estimated distribution of refillable coefficients) and $\rho \geq 0$. The value ρ affects the degree of heterogeneity in the refillable coefficient and affects how customers sort into the different formats. At $\rho = 0$, all customers dislike refillables the same, but this changes as ρ increases. Price-insensitive customers will dislike refillables relatively more as ρ increases, implying that the higher the value of ρ , the more perfect the customer sorting that the retailer can achieve using the refillable segment. Note that at $\rho = 1$, the counterfactual distribution is equal to the estimated distribution of preferences.

Figure 6 shows the results of the profit analysis for different counterfactual distributions (indexed by ρ) for the Diet Coke products.²² The figure shows that profits with both formats $(\pi(\rho)/\pi(1))$ are uniformly greater than the profits without the refillables $(\pi^{\text{No Ref}}/\pi(1))$. This result captures that, regardless of the value of ρ , introducing refillables expands quantity and affects profits positively as refillables become a better fit for some customers. However, the difference between $\pi(\rho)/\pi(1)$ and $\pi^{\text{No Ref}}/\pi(1)$ varies importantly when changing ρ . For low values of ρ , price-insensitive customers are not particularly deterred from purchasing refillables, which prevents the retailer from targeting low prices to price-sensitive customers without attracting price-insensitive customers as well. As the value of ρ increases, fewer price-insensitive customers are willing to bear the inconvenience cost of purchasing

²² Similar figures for Coke and Coke Zero show qualitatively identical results and are available upon request.




FIGURE 6. MEASURING THE CHANGE IN PROFITS AS A FUNCTION OF THE DEGREE OF HETEROGENEITY IN THE TASTE FOR REFILLABLES: DIET COKE PRODUCTS

Notes: ρ is the coefficient on expression $\gamma_i^{\text{Counterfactual}}(\rho) = \underline{\gamma} + \rho(\gamma_i - \underline{\gamma})$, where $\underline{\gamma} = \min_{i \in I} \gamma_i$. $\pi^{\text{No ref}}/\pi(1)$ is the value of profits when refillables are removed as a share of profits when both formats are available. Both $\pi^{\text{No ref}}$ and $\pi(1)$ were computed using the estimated preferences (i.e., $\rho = 1$). $\pi(\rho)/\pi(1)$ are profits (including both formats) for counterfactual preferences $(\alpha_i, \gamma_i(\rho))$ as a share of the profits when $\rho = 1$.

in the refillable format, which allows the retailer to target the lower prices to the price-sensitive customers.

The results in Figure 6 suggest that the gains of introducing the refillable products into the market are magnified when the refillables allow the seller to target lower prices to price-sensitive customers. With the class of counterfactual customer heterogeneity distributions given in (8), one observes that the profit gains of introducing refillables tend to disappear as the customer heterogeneity that enables customer sorting fades away.

V. Conclusion

In this paper I study a market where customers can purchase soda in a refillable or nonrefillable bottle. Purchasing refillables allows access to lower prices but at the cost of an inconvenience, as customers have to return a refillable bottle before each purchase.

I argue that the profitability of this market segmentation depends on how the price sensitivity and taste for purchasing in the refillable format of customers are jointly distributed. Using a discrete choice model, I estimate this joint distribution and find a negative correlation between price coefficient and the taste for purchasing in the refillable format. This relationship allows the retailer to lower the prices of refillables to target lower prices to price sensitive customers without attracting all price-insensitive customers.

When evaluating the effects of market segmentation, I find that profits and average customer welfare would decrease by 4.21 and 12.61 percent, respectively, if the refillable products were removed from the market. This result reflects that the more price-sensitive customers are forced to exit the market as a consequence of the

cheaper refillable products no longer being available. I also find that the gains of serving customers with both formats would be lower if the price-insensitive customers were not particularly deterred from purchasing refillables. This result suggests that the value of serving customers with both formats depends on the nature of customer heterogeneity, as it dictates whether the seller can use the refillable segment to price discriminate.

REFERENCES

- Berry, Steven, and Ariel Pakes.** 2007. "The Pure Characteristics Demand Model." *International Economic Review* 48 (4): 1193–1225.
- Borenstein, Severin.** 1991. "Selling Costs and Switching Costs: Explaining Retail Gasoline Margins." *RAND Journal of Economics* 22 (3): 354–69.
- Borenstein, Severin, and Nancy L. Rose.** 1994. "Competition and Price Dispersion in the U.S. Airline Industry." *Journal of Political Economy* 102 (4): 653–83.
- Busse, Meghan, and Marc Rysman.** 2005. "Competition and Price Discrimination in Yellow Pages Advertising." *RAND Journal of Economics* 36 (2): 378–90.
- Clerides, Sofronis K.** 2004. "Price Discrimination with Differentiated Products: Definition and Identification." *Economic Inquiry* 42 (3): 402–12.
- Deneckere, Raymond J., and R. Preston McAfee.** 1996. "Damaged Goods." *Journal of Economics and Management Strategy* 5 (2): 149–74.
- Eizenberg, Along, and Alberto Salvo.** 2012. "Grab them Before they Go Generic: Habit Formation and the Emerging Middle Class." <http://cep.lse.ac.uk/seminarpapers/11-10-12-AE.pdf>.
- Goldberg, Pinelopi Koujianou.** 1996. "Dealer Price Discrimination in New Car Purchases: Evidence from the Consumer Expenditure Survey." *Journal of Political Economy* 104 (3): 622–54.
- Hendel, Igal, and Aviv Nevo.** 2006. "Measuring the Implications of Sales and Consumer Inventory." *Econometrica* 74 (6): 1637–73.
- Hendel, Igal, and Aviv Nevo.** 2013. "Intertemporal Price Discrimination in Storable Goods Markets." *American Economic Review* 103 (7): 2722–51.
- Hviid, Morten, and Greg Shaffer.** 1999. "Hassle Costs: The Achilles' Heel of Price-Matching Guarantees." *Journal of Economics and Management Strategy* 8 (4): 489–521.
- Kaplan, Greg, and Guido Menzio.** 2014. "The Morphology of Price Dispersion." http://scholar.princeton.edu/sites/default/files/gkaplan/files/kaplan_menzio_ier_2014.pdf.
- Leslie, Phillip.** 2004. "Price Discrimination in Broadway Theater." *RAND Journal of Economics* 35 (3): 520–41.
- Lesser, William, and Ananth Madhavan.** 1987. "Economic Impacts of a National Deposit Law: Cost Estimates and Policy Questions." *Journal of Consumer Affairs* 21 (1): 122–37.
- Marshall, Guillermo.** 2015. "Hassle Costs and Price Discrimination: An Empirical Welfare Analysis: Dataset." *American Economic Journal: Applied Economics*. <http://dx.doi.org.10.1257/app.20130046>.
- McFadden, Daniel.** 1981. "Econometric Models of Probabilistic Choice." In *Structural Analysis of Discrete Data with Econometric Applications*, edited by Charles F. Manski and Daniel L. McFadden, 198–272. Cambridge: MIT Press.
- McManus, Brian.** 2007. "Nonlinear pricing in an oligopoly market: The case of specialty coffee." *RAND Journal of Economics* 38 (2): 512–32.
- Mortimer, Julie Holland.** 2007. "Price Discrimination, Copyright Law, and Technological Innovation: Evidence from the Introduction of DVDs." *Quarterly Journal of Economics* 122 (3): 1307–50.
- Mussa, Michael, and Sherwin Rosen.** 1978. "Monopoly and Product Quality." *Journal of Economic Theory* 18 (2): 301–17.
- Nevo, Aviv.** 2011. "Empirical Models of Consumer Behavior." *Annual Review of Economics* 3: 51–75.
- Nevo, Aviv, and Catherine Wolfram.** 2002. "Why do manufacturers issue coupons? An empirical analysis of breakfast cereals." *RAND Journal of Economics* 33 (2): 319–39.
- Pakes, Ariel, and David Pollard.** 1989. "Simulation and the Estimation of Optimization Estimators." *Econometrica* 57 (5): 1027–57.
- Petrin, Amil, and Kenneth Train.** 2010. "A Control Function Approach to Endogeneity in Consumer Choice Models." *Journal of Marketing Research* 47 (1): 3–13.
- Porter, Richard C.** 1978. "A Social Benefit-Cost Analysis of Mandatory Deposits on Beverage Containers." *Journal of Environmental Economics and Management* 5 (4): 351–75.

- Shepard, Andrea.** 1991. "Price Discrimination and Retail Configuration." *Journal of Political Economy* 99 (1): 30–53.
- Song, Minjae.** 2012. "A Hybrid Discrete Choice Model of Differentiated Product Demand with An Application to Personal Computers." <http://www.simon.rochester.edu/fac/MSONG/papers/Song-hybrid.pdf>.
- Tirole, Jean.** 1988. *The Theory of Industrial Organization*. Cambridge: MIT Press.
- Varian, Hal R.** 1985. "Price Discrimination and Social Welfare." *American Economic Review* 75 (4): 870–75.
- Varian, Hal R.** 1989. "Price Discrimination." In *Handbook of Industrial Organization*, Vol. 1, edited by Richard Schmalensee and Robert Willig, 597–654. Amsterdam: North-Holland.
- Villas-Boas, Sofia Berto.** 2009. "An empirical investigation of the welfare effects of banning wholesale price discrimination." *RAND Journal of Economics* 40 (1): 20–46.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.