To Discriminate or Differentiate?

Evidence From a Supermarket Chain

Charles Taragin*

November 27, 2006

Abstract

Retailers with market power can extract consumer surplus by introducing new products (horizontal differentiation) or by offering existing products in different sizes (price discrimination). Both strategies have their advantages; spatial Hotelling models suggest that a retailer can increase its profits by adding new horizontally differentiated products and then raising the price on all of its products. Models of second-degree price discrimination also suggest that retailers can raise profits by offering a menu of price\size alternatives and allowing customers to sort themselves into the alternative that most suits them.

While both horizontal differentiation and price discrimination provide retailers with distinct opportunities to increase profits, they are not necessarily complementary strategies. I employ a model of second-degree price discrimination to show that the introduction of a horizontally differentiated product causes the prices and profits on an existing product's menu alternatives to decrease. This result implies that retailers must sacrifice some of their ability to price discriminate in order to horizontally differentiate.

Using detailed, store-level data from the Dominick's Fine Foods supermarket chain, I investigate whether stores within this chain trade the return from offering an existing sports drink in different container sizes (price discrimination) for the return from introducing new products (horizontal differentiation).

I find evidence suggesting that such an exchange does in fact occur.

*These results are preliminary, and I welcome your comments. I wish to thank Patrick Bajari, Francine LaFontaine, Kai-Uwe Kuhn, Justin McCrary, Ron Bozekowski and Raphael Thomadsen for all of their help. I would also like to thank the James M. Kilts Center, GSB, University of Chicago for providing access to the Dominick's scanner data.

1 Introduction

Retailers with market power can extract consumer surplus by introducing new products (horizontal differentiation) or by offering existing products in different sizes (price discrimination). Both strategies have their advantages; spatial Hotelling models suggest that a retailer can increase its profits by adding new horizontally differentiated products and then raising the price on all of its products. Models of second-degree price discrimination also suggest that retailers can raise profits by offering a menu of price\size alternatives and allowing customers to select the alternative that most suits them.

While both horizontal differentiation and price discrimination provide retailers with distinct opportunities to increase profits, they are not necessarily complementary strategies. I employ a model of second-degree price discrimination to show that the introduction of a horizontally differentiated product causes the prices and profits on an existing product's menu alternatives to decrease. This result implies that retailers must sacrifice some of their ability to price discriminate in order to horizontally differentiate.

Using detailed, store-level data from the Dominick's Fine Foods supermarket chain, I then investigate whether stores within this chain trade the return from offering an existing sports drink in different container sizes (price discrimination) for the return from introducing new products (horizontal differentiation). Sports drinks are an ideal choice for studying this trade-off; unlike other products sold in Dominick's stores, sports drinks are almost completely characterized by their size, manufacturer, and flavor, all of which are observed in the data. Moreover, since the number of sports drinks sold by Dominick's stores increases over the course of the sample, the data contain a number of opportunities to examine the effect of new product entry.

My investigation into the relationship between differentiation and discrimination is complicated by the fact that stores choose which sports drinks to stock as well as how to price these products. These choices are contingent on a number of store-specific factors not observed in the data, including competition between neighboring grocers, the vertical relationship between stores and sports drink distributors, and whether a store chooses to prominently display a sports drink product.

I use the delay in a store's receipt (removal) of a product as part of a differences-in-differences strategy to distinguish the effect of new product entry from the effects of these unobserved factors. Provided that the order in which stores receive (remove) products is uncorrelated with unobserved factors affecting a store's pricing policy, differences-in-differences allows me to identify the causal effect of product entry (exit) on an existing product's price schedule.

The above strategy assumes that conditional on observables, all existing products will be affected by another product's entry or exit in a similar fashion. Theory suggests, however, that the degree to which products substitute for one another plays an important role in this process. Products that are not close substitutes for a new product should experience little distortion in their price schedules and menu offerings when a new product enters, while products that are close substitutes should experience substantial distortions. As a result, the differences-in-differences strategy should understate the causal effect of product entry. I remedy this by first using Bayesian methods to construct own-price elasticities for the weeks surrounding a product release (retirement), and then incorporating these elasticities into the differences-in-differences strategy.

Few empirical researchers have examined how multi-product retailers with market power exchange price-discrimination for horizontal differentiation. Draganska and Jain (2006) use a structural model to analyze whether consumers value horizontally differentiated yogurt flavors more than vertically differentiated yogurt product lines. Having established that consumers value line attributes more than flavor attributes, Draganska and Jain go on to determine whether manufacturers use product lines to price discriminate. While this analysis represents the first step in examining how yogurt manufacturers might trade flavors for products, Draganska and Jain do not explicitly examine this exchange.

A number of empirical studies have also examined a related problem, the relationship between competition and price discrimination. Busse and Rysman (2005) find that competition increases the curvature in the price schedule of yellow pages advertisements, while Seim and Viard (2004) find that increased competition leads to a proliferation of wireless calling plans. Likewise, Borzekowski et al. (2006) find that increased competition amongst direct mailers is associated with an increased propensity to price discriminate.

The main difference between these studies and the one that I propose is that Dominick's supermarkets are able to choose which horizontally differentiated products to stock as well as how to price these products, while firms in these studies choose their prices *conditional* on their competitors' product offerings and price schedules. This distinction implies that Dominick's supermarkets can better manage the trade-off between differentiation and discrimination than firms facing competition.

2 A Model of Price Discrimination

I use a variant of the canonical second-degree price discrimination problem proposed by Itoh (1983) to explore the relationship between horizontal differentiation and price discrimination. Itoh's model describes a single-product monopolist who is able to engage in nonlinear pricing by offering a menu of prices and quantities[1] to consumers with ranked, heterogeneous tastes for the monopolist's product. Itoh (and others) show that even if a monopolist only knows the distribution of consumer tastes, it can induce consumers to reveal their tastes by appropriately choosing a sequence of price-quantity menus.

While this model explicitly deals with only one product, in fact two products are present; the good produced by the monopolist and the "outside good" whose value must be accounted for by the monopolist to ensure that consumers are willing to purchase its product. The outside good plays the part of the next best alternative to the monopolist's product, and can be thought of as the value of a horizontally differentiated product.

Unfortunately, Itoh normalizes the value of the outside good to 0, making it impossible to examine how differentiation and discrimination trade-off with one another. Below, I present Itoh's model with this assumption relaxed and demonstrate how the introduction of a new product (which I interpret as an increase in the value of the outside option) affects the monopolist's ability to price discriminate.

[1] Itoh's model deals with quality and not quantity-based price discrimination. Maskin and Riley (1984) demonstrate that these two types of discrimination are equivalent.

2.1 Consumers

Itoh (1983) describes a single-product monopolist who offers consumers M versions of its product, where the version menu is described by the sequence of price quantity pairs $A = \{(p_m, q_m)\}_{m=1}^{M}$, with $q_1 > q_2 > \ldots > q_M$. Each consumer faces the option of either purchasing one of the M versions in A or not purchasing the product. The utility of the ith consumer is described by

$$U_{im} = \begin{cases} \theta q_m - p_m & \text{if } i \text{ purchases } \exists m \in A \ ; \\ \underline{u} & \text{otherwise.} \end{cases}$$

Consumers are endowed with some $\theta \in [\underline{\theta}, \bar{\theta}]$ that dictates how they value higher quantities of the monopolist's product; consumers endowed with larger θs value additional quantity more than consumers with smaller θs. Any consumer can choose not to purchase from the monopolist's menu and instead earn utility $\underline{u}$.

Both θ and $\underline{u}$ play an important role in my analysis. θ represents the vertical dimension over which the monopolist price discriminates, while $\underline{u}$ represents the value of a horizontally differentiated product.

Itoh shows that even if the monopolist only knows the distribution of consumer-types, the monopolist can still prevent consumers with higher θs from choosing product versions meant for consumers with lower θs. To accomplish this, Itoh notes that a consumer-type will purchase version m only if the consumer prefers version m to not purchasing any version of the product and doesn't prefer version m to $m+1$. These conditions can be written as

$$\begin{aligned} \theta q_m - p_m &\geq \underline{u} \\ \theta q_m - p_m &\geq \theta q_n - p_n, n \in [m+1, M] \end{aligned}$$

where the first condition is referred to as the individual rationality constraint, while the second is referred to as the incentive compatibility constraint.

Itoh defines θ_m as the consumer-type indifferent between purchasing menu versions m and $m+1$ and shows that the individual rationality constraint only applies to type θ_M consumers, while the incentive compatibility constraints only apply to adjacent menu items. Hence

$$\begin{aligned} \theta_M q_M - p_M &= \underline{u} \\ \theta_m q_m - p_m &= \theta_m q_{m+1} - p_{m+1}, \forall m < M \end{aligned}$$

which implies that the market demand for menu option m is given by

$$\int_{\theta_m}^{\theta_{m-1}} f(\theta) d\theta.$$

with $\theta_0 = \bar{\theta}$, $\theta_m = \frac{p_m - p_{m+1}}{q_m - q_{m+1}}, \forall m < M$ and $\theta_M = \frac{\underline{u} + p_M}{q_M}$

Solving the second equation for p_{m+1} and recursively substituting yields

$$\begin{aligned} p_m &= \sum_{j=m}^{M-1} \Delta p_j + p_M \\ p_M &= \theta_M q_M - \underline{u}. \end{aligned} \tag{1}$$

where $\Delta p_m = p_m - p_{m+1}$

Equation 1 is another way of describing the well known result that consumers of a particular type only want to deviate to the lower quantity menu options. This result will also prove useful in determining how the value of the outside option affects the price schedule.

2.2 The Monopolist

Suppose that the monopolist is selling M versions of its product. Then the monopolist must choose a sequence of prices to solve

$$\max_{\{p_m\}_{m=1}^{M}} \sum_{m=1}^{M} \int_{\theta_m}^{\theta_{m-1}} (p_m - C(q_m)) f(\theta) d\theta$$

which Itoh rewrites as

$$\max_{p_M, \{\Delta p_m\}_{m=1}^{M-1}} \int_{\theta_M}^{\bar{\theta}} (p_M - C(q_M)) f(\theta) d\theta + \sum_{m=1}^{M-1} (\Delta p_m - (C(q_m) - C(q_{m+1})) \int_{\theta_m}^{\bar{\theta}} f(\theta) d\theta$$

where $C(q_m)$ is the cost associated with producing quality level q_m .

Differentiating with respect to $p_M, \{\Delta p_m\}_{m=1}^{M-1}$ yields

$$\begin{aligned} p_M &= (q_m / f(\theta_M)) \cdot \int_{\theta_M}^{\bar{\theta}} f(\theta) d\theta + C(q_M) & (2) \\ \Delta p_m &= (q_m - q_{m+1} / f(\theta_m)) \cdot \int_{\theta_m}^{\bar{\theta}} f(\theta) d\theta + (C(q_m) - C(q_{m+1})), \forall m < M & (3) \end{aligned}$$

Equation 2 implies that the price of the smallest menu (p_M) is independent of the prices of larger menu items, while equation 3 implies that the price of menu m is independent of all prices except for the price of the next smallest quantity . As such, only p_M effects the total demand for the monopolist's product, while only the price of the next smallest quantity directly effects the demand for a particular menu item.

2.3 Some Theoretical Predictions

Itoh uses the above model to show the following

Proposition 1. *A menu item's price only depends on the prices of smaller menu items.*

Proposition 2. *The demand for the monopolist's product only depends upon the price of the smallest menu item.*

Proposition 3. *Introducing or removing a menu item changes the prices of larger menu items.*

To see how changing the value of the outside option changes menu prices, I differentiate the first order conditions with respect to $\underline{u}$ to obtain

$$\frac{dp_M}{d\underline{u}} = -\frac{f'(\theta_M)}{f^2(\theta_M)} \cdot \int_{\theta_M}^{\bar{\theta}} f(\theta)d\theta - 1 \tag{4}$$

$$\frac{d\Delta p_m}{d\underline{u}} = 0, \forall m < M \tag{5}$$

Combining this result with equation 1 yields

$$\begin{aligned} \frac{dp_m}{d\underline{u}} &= \frac{dp_M}{d\underline{u}} \\ &= -\frac{f'(\theta_M)}{f^2(\theta_M)} \cdot \int_{\theta_M}^{\bar{\theta}} f(\theta)d\theta - 1, \forall m < M \end{aligned} \tag{6}$$

which is negative if $f'(\theta_M) > \frac{-f^2(\theta_M)}{\int_{\theta_M}^{\bar{\theta}} f(\theta)d\theta}$. Equation 6 demonstrates that

Proposition 4. *Changes in the value of the outside option change the prices of **all** menu items by the same amount.*

Moreover, since $\frac{d\theta_M}{d\underline{u}} = \frac{1}{q_m} > 0$, by Proposition 2, increasing the value of the outside option shrinks the interval $[\theta_M, \bar{\theta}]$, decreasing the total demand for the monopolist's product. When combined with Proposition 4, this implies

Proposition 5. *If $f'(\theta_M) > \frac{-f^2(\theta_M)}{\int_{\theta_M}^{\bar{\theta}} f(\theta)d\theta}$, then increasing the value of the outside option never increases the monopolist's profits from an existing menu item.*

2.4 Discussion

Proposition 5 highlights one difficulty with using this model to ascertain how horizontal differentiation and price discrimination relate to one another. First, suppose a monopolist introduces a new product at a single price that is a close substitute to an existing product, and that the monopolist continues to price discriminate over the existing product. According to the model described above, only the low quality types will have any incentive to switch to the new product, and all demand for the new product will either be from low types or from customers who were not market participants.

Rochet and Stole (2002) remedy this odd implication by assuming that consumers are endowed with two random types,the first which dictates their preference for quality, and the second which dictates how they value the outside good. Rochet and Stole describe how this second random variable adds horizontal differentiation to Itoh 's model of vertical differentiation, eliminating his result that the individual rationality constraint only binds for the lowest type. With this modification, a uniform price increase across all menus causes all types to decrease demand. Similarly, the introduction of a new, closely substitutable product should increase the option value of the outside good, decreasing the demand for all menus. Unfortunately, Rochet and Stole do not offer concrete predictions about how the introduction of a new product *by the same firm* affects the price schedule of an existing product.

3 Data

I use data on supermarkets in the Dominick's supermarket chain obtained from the University of Chicago's Kilts' Marketing Center to test some of the above predictions. Kilts maintains a database containing weekly sales, retail price, wholesale price, and display information for many of the SKUs sold at each of the 93 supermarkets in the Dominick's Fine Foods chain over a 7 year period beginning in 1989. SKU stands for Stock Keeping Unit, or the level at which a unit of sale is identified in the data. Here I can identify an SKU by its manufacturer, flavor and size. I define a product as the collection of SKUs with the same manufacturer and flavor.

The Kilt's data also houses a database of store-level characteristics, including data from the Census blocks surrounding each store, as well as information about each stores' weekly revenue and attendance.

3.1 Sports Drinks

These data cover an important period in the sports drink market. Over the course of the sample, Gatorade's twenty-year monopoly faced two major challenges, first by Coca-Cola's Powerade product line in 1992, followed a year later by Pepsi's All-Sport line. Table 1 displays the total sales and U.S market shares of each firm from 1988-1997. This table demonstrates that the sports drink market underwent a major expansion over the course of the sample; total sales doubled from 1988-1994. Moreover, despite its best efforts, Gatorade was unable to prevent either All-Sport or Powerade from eroding its position; its market share fell from 82.4%, the year that All-Sport first entered to 73.1% by 1997.

Figure 1 demonstrates that competition amongst these manufacturers unfolded somewhat differently across Dominick's stores. Unlike the national market, where Powerade was the first to compete against Gatorade, All-Sport was the first to compete against Gatorade in Dominick's stores, entering the market in the 183rd week of the sample. Powerade's entry was postponed until the 242nd week, more than a year after All-Sport's entry. Apparently, being the earlier entrant was advantageous to All-Sport; unlike the national market where their share lagged Powerade's, All-Sport market share dominated Powerade's for most of the sample. Figure 1 also indicates that Dominick's stores differ from the national market in one important way; although Gatorade is clearly affected by both All-Sport and Gatorade's presence, the figure indicates that by the end of the sample, Gatorade had regained virtually the entire market.

Table 2 lists the 68 sport drink SKUs that are present over the course of the sample[2] by manufacturer, flavor, size, and its first and last appearance in the sample. Notice how Gatorade's products dominate the data; Gatorade's has 46 SKUs in the sample, almost 3 times All-Sport's and 4 times Powerade's. Gatorade has more flavors than either Gatorade or All-Sport (14, compared to 6 and 7) and more sizes per flavor (a median of 4 compared to 3.5 and 2). Gatorade is also responsible for virtually all of the new flavors present in Dominick's stores; 6 of the 9 new flavors that entered Dominick's were produced by Gatorade (All-Sport had one new flavor, Powerade had two new flavors).

[2] I ommit 6 SKUs from the sample; 5 of these SKUs have no identifiable flavor information and the 6th is omitted because it is the sole SKU sold in powder form.

3.2 Dominick's Stores

Dominick's corporate office typically chooses which products its stores carry and which prices they charge. Chintagunta et al. (2003) describe how Dominick's uses zone pricing as a form of third-degree price discrimination. According to Chintagunta et al., Dominick's zone policy evolved over time, growing from 3 zones in 1990 to 16 pricing zones by 1992. While stores within a zone rarely sell a product for the same price, prices are less dispersed within a zone than across zone.

Chintagunta et al.'s finding that that Dominick's engages in a form of third degree price discrimination suggests that Dominick's stores possess some degree of market power. Dominick's ability to price discriminate is supported by Hoch et al. (1995), who find that that while store-level category [3] price elasticities are sensitive to the proximity of competitors, competitive effects are relatively unimportant when compared to measures of store market power.

In addition to setting prices, Dominick's corporate office also decides which products individual stores will carry. The corporate office maintains a team of corporate buyers who maintain a particular product category. When a new product is released, "roll-out" teams go from store to store to install a new product in its category. These buyers also determine when to remove a product from their shelves.

Store managers also play a role in this process. They pass on consumer requests for a particular product to the corporate office, and if corporate decides to no longer carry a particular product, can negotiate directly with a product's manufacturer to obtain a product.[4]

Figure 2 succinctly displays how the sports drink inventory of each of Dominick's stores differ from one another. I construct figure 2(a) by first calculating the fraction of all SKUs that each store sells in a particular week, and then plotting that fraction for stores at the 75th (dashed), 50th (solid), and 25th (dashed) percentiles. The leftmost horizontal line marks All-Sport's entry into the chain, while the rightmost line marks Powerade's entry. Figure 2(b) is similar to figure 2(a) except that it displays the fraction of all products sold by each store during a given week.

Taken together, these graphs suggest that prior to All-Sport's introduction, each store carried

[3] categories are collections of products. Examples include cookies, juice, sodas, soups, and cigarettes. Sports drinks are listed in both the bottled juice and carbonated sodas categories.

[4] I obtained this institutional detail from an interview with a Dominick's store manager.

all of Gatorade's flavors but that virtually no store stocked all the sizes that each product was offered in. The increase in store shares that occurs before All-Sport's entry could either be a result of each store carrying more sizes or Dominick's corporate discontinuing some of the sizes. Table 2 indicates that during this period, Dominick's stopped selling four SKUs but began selling 6 new SKUs, which suggests that stores began carrying more sizes.

After All-Sport's release, these graphs show that stores began carrying a smaller fraction of all SKUs as well as a smaller fraction of all products. An explanation again comes from table 2, which reveals that all of All-Sport's and Powerade's flavors at entry duplicated Gatorade's flavors, and that virtually all of Gatorade's products are sold throughout the sample. These observations suggest that stores chose to conserve scarce shelf-space by not offering duplicate flavors.

3.3 Do Supermarkets Price Discriminate?

Having established that stores in the Dominick's chain horizontally differentiate, I must now determine whether or not stores price discriminate. Theory predicts that a price-discriminating monopolist's pricing schedule will be concave in size. To test this prediction, I regress

$$\begin{aligned} \log\left(\textit{retail price}_{ijst}\right) &= \beta_0 + \beta_1 \textit{wholesale price}_{ijst} + \beta_2 \textit{bundle}_i + && (7) \\ & \quad \delta_v + \alpha_s + \gamma_j + \phi_m + \omega_t + \epsilon_{ijst} && (8) \end{aligned}$$

where γ_j is the coefficient on the product j dummy to which SKU i belongs, and ϕ_m, α_s, are the coefficients on manufacturer, store, and week dummies. δ_v are the coefficients on a series of volume dummies that measure the percentage change in price due to the number of fluid ounces in SKU i , relative to the 16 ounce size. $\textit{bundle}_i$ is an indicator variable equal to 1 if an SKU is a bundle of other SKUs and is included to distinguish the 16 ounce 4-pack from the 64 ounce size.

Equation 7 uses price variation across stores, weeks, and products to estimate δ_v. Normally, δ_v would not identify price discrimination because price variation could be due differences in the marginal cost of producing larger sizes, and I would need to implement a strategy similar to Shepard

(1991) or Cohen (2004) to identify price discrimination. Fortunately, the data include a measure of SKU wholesale prices, making these strategies unnecessary.

Figure 3 displays δ_v, confirming the theoretical prediction[5]; retail prices are indeed concave in container size.

4 Estimation Strategies

Having established that stores both differentiate and discriminate, I use a variant of the differences-in-differences (DIFF) to examine how the introduction or removal of a horizontally differentiated product affects the prices of an existing product's menu items. Proposition 5 states that introducing (removing) a horizontally differentiated good changes prices and profits on all of the existing product's menu items. As such, I use the price markup before and after a product is released (or removed) to determine whether stores are exchanging their ability to price discriminate for the opportunity to differentiate.

Unfortunately, the DIFF strategy assumes that conditional on an SKU's observables, the addition or subtraction of a product always affects existing SKUs in the same way. For example, suppose that a store selling only the lemon-lime flavor introduces the tropical fruit flavor in week 82 and the grape flavor in week 126. The DIFF strategy assumes that on average, the introduction of the grape flavor will have a similar effect on the lemon-lime and tropical fruit flavors as the introduction of the tropical fruit flavor had on the lemon-lime flavor.

To see why this might prove false, suppose that consumers prefer to not purchase any sports drink to purchasing the tropical fruit flavor but always prefer grape to both lemon-lime and tropical fruit. Then the introduction of the tropical fruit flavor, does not change the value of the outside option and should therefore not affect the price schedule of lemon-lime SKUs. On the other hand, since consumers always prefer grape to both lemon-lime and tropical fruit, the introduction of the grape flavor should increase the value of the outside option, changing the price schedule for both lemon-lime and tropical fruit SKUs. In this instance, the DIFF strategy mistakenly averages the outcome of the two new flavor introductions together, understating the effect of a close substitute.

[5] All size dummies are statistically significant at the 99% level.

I remedy this problem by structurally estimating an SKU's demand curve and conditioning the DIFF strategy on the estimated own-price elasticities. Doing so allows each SKU to respond differently to a product entry or exit, enabling the DIFF strategy to distinguish products that affect the value of the outside option from products that do not.

The estimation strategy described above will only capture the trade-off between discrimination and differentiation if the following assumptions hold. First, I assume that stores are embedded in distinct markets. This assumption implies that consumers find it too costly to travel from one supermarket to another searching for deals on sports drinks. If this assumption fails, then competition amongst subsets of stores, and not the introduction of new products could affect the tests described above. For similar reasons, I assume that competition between Dominicks and other supermarket chains (like Jewel-Osco) occurs at the chain rather than at the store level.

4.1 Estimating Differences in Differences

The main objective of my DIFF analysis is to estimate the causal effect of product's introduction (removal) on a supermarket's ability to price discriminate over one of its existing (remaining) products, *conditional on that supermarket eventually adopting the new product.* Conditioning on adoption ensures that my estimates will remain unbiased even if my empirical specification does not properly account for why some supermarkets choose to never adopt (remove) a product.

Why use DIFF? A simple comparison of retail price markups before and after a product release (removal) should indicate how firms trade off discrimination and differentiation. Unfortunately, this comparison will only capture the effect of interest if no other factors are driving the store's decision to introduce (retire) the product. If, say, stores add (remove) products because demand for sports drinks increases (decreases) then the resulting change in the outcome of interest will capture both the change due to the demand increase (decrease) as well as the effect on price discrimination.

DIFF solves this problem by identifying a control group; in this case, stores that experience the demand increase but are delayed in receiving the new flavor. For this group of stores, the change in the outcome is due entirely to the change in demand. DIFF is constructed by subtracting the change in the outcome due exclusively to the change in demand from the simple difference described above, thereby isolating the effect of price discrimination.

4.1.1 Store Adoption (and Removals)

One question remains; why would some stores receive (remove) products before others? This is not a frivolous question: if the stores who receive (remove) the product relatively later (the control group) also experience the demand increase later, then the control group no longer measures the change in the outcome due only to the change in demand. As a result, DIFF captures both the change in the outcome due to the change in demand as well as the change due to the new product's introduction.

I maintain that conditional on the chain adopting (removing) an SKU, variation in store adoption times depends on logistical issues that are largely unrelated to the conditions in any particular store. In section 3, I argued that Dominick's "roll out" teams were responsible for setting up (and removing) SKUs across stores, which if true would imply that my assertion is correct.

Table 6 provides evidence supporting my claim. There, I identify the first week that a new product appears in one of Dominick's stores and then examine the fraction of eventual adopters who adopt that product for 10 weeks after the product's first appearance. This table reveals that virtually every store that eventually adopts a product does so by the 10th week following the product's introduction, with (on average) 3/4 of all stores adopting a product after the 3rd week.

Table 6 also demonstrates that much of the variation in store adoption times is manufacturer related. Many of Powerade's and Gatorade's products are almost universally adopted after the 3rd week, while All-Sport's products take until the 8th week to achieve similar penetration. Powerade's Mountain Blast and Tidal Burst flavors are the exception to this pattern; both enter stores at markedly slower rates than other Powerade products.

Taken together, these facts indicate that i) the Dominick's chain and not individual stores choose whether to stock a product and ii) Manufacturers differ in their ability to distribute new products to stores.

Exiting products tell a similar story. Table 6 displays the fraction of retirees who discontinue a product 10 weeks before the final retiree discontinues the product. The fraction of retirees still selling the product remains relatively constant until 4 weeks before the final store retires, when it decreases exponentially. This suggests that like new products, Dominick's chain sets target dates for the removal products from certain stores.

The main difference between the entry and exit is that some stores apparently choose to discontinue products well before the Dominick's chain requires them to. For instance, table 6 indicates that 10 weeks before the final retirement, roughly half of all retirees had already stopped carrying some of All-Sport's products. For these retirees, it seems likely that store-specific events and not chain policy lead retirees to stop selling the product. If these events are corellated with a store's pricing strategy then my estimation strategy will not identify the causal effect of entry (exit).

4.1.2 Specification

I implement a version of DIFF using the following specification

$$\log\left(retail\ price_{ist}\right) = \beta_0 + \varphi wholesale\ price_{ist} + \kappa_1 samefirm_{ist} + \kappa_2 sameproduct_{ist} + \kappa_3 sameflavor_{ist} + \sum_{l=-4}^{4} \delta_l treat_{lst} + \beta_t + \alpha_s + \gamma_i + \phi_m + \epsilon_{ist} \quad (9)$$

$samefirm_{ist}$, $sameproduct_{ist}$, and $sameflavor_{ist}$ are dummy variables indicating whether the entering SKU is either produced by the same manufacturer, part of the same product, or similarly flavored. These dummies are included to account for differences between entering (exiting) products.

Following Jacobson et al. (1993), I capture the causal effect of product entry and exit by defining the following dummies

$$treat_{lst} = \begin{cases} 1 & \text{if a product enters (exits) store } s\ l \text{ periods before week } t; \\ 0 & \text{otherwise.} \end{cases}$$

These dummies play an integral role in my analysis. To see how, recall that SKUs enter and leave the Dominick's chain throughout the sample, making these different events difficult to compare. The relative time dummies $\delta = (\delta_{-4}, \ldots, \delta_4)$ ease this comparison by normalizing entry (exit) to a common week 0.

δ serves two purposes. Absent any other factors that influence price, δ represents the trend in SKU i's retail price in the weeks surrounding a product release (retirement). In particular, $\delta_0 - \delta_{-1}$ captures the causal effect of new product entry (exit). The remaining elements of δ serve as diagnostic tool; they highlight retail price trends that are not captured by other regression covariates.

Demand shocks, including holidays, sporting events, and weather could obscure the effect of product entry (exit). I control for these phenomena by including week fixed effects (β_t). I also control for differences in consumers across stores and changes in those consumers purchasing habits by including store-level fixed effects (α_s). These store level effects control for any observed store-level heterogeneity that is constant over the sample. For similar reasons, I include SKU-level fixed effects (γ_i) and manufacturer fixed effects (ϕ_m) in my analysis.

I choose a four week window around a product's entry (exit) to minimize the overlap between adjacent events and estimate equation 9 using ordinary least squares with robust standard errors. I present my results in section 5.

4.2 Estimating Demand

As discussed earlier, the DIFF strategy presented above will understate the effect of product entry (exit) on existing products if entrants (retirees) affect some products differently than others.

To control for these effects, I estimate an SKU's demand using a random coefficient multinomial logit. Numerous studies employ this demand model because of its economy; as Berry (1994) points out, even simple linear demand system with J products require the econometrician to estimate J^2 parameters. In my case, this problem is compounded since I expect these parameter estimates to differ before and after the products release. The multinomial logit remedies this problem by translating the econometrician's problem into characteristic space which has a smaller dimension than product space.

This economy comes at a cost. Controlling for endogeneity when using the multinomial logit's is difficult. While Berry et al.'s two-step estimator does account for certain forms of endogeneity, their algorithm may not converge if the econometrician happens to choose poor starting values. To remedy this, I adapt Bayesian methods from Rossi et al. (2005) to estimate a random coefficient multinomial logit model over grouped data. The Bayesian method has one great advantage over other methods; because it returns the joint distribution of model coefficients in the sample rather than just the distribution's moments, it is more robust to the choice of initial values.

To simplify the estimation, I aggregate stores up to the pricing zone level and then estimate consumer utility by zone and week for 4 weeks before and after a new product is introduced. Aggregating stores to the pricing zone is reasonable since Dominick's pricing zones group stores with similar consumer populations together.

4.2.1 Shopper Behavior

Suppose that a shopper in pricing zone z during week t can choose to consume one of J_{zt} sports drink SKUs sold in zone z during week t. The shopper earns (indirect) utility

$$V_{zjt} = \alpha_{zt}(income_{zt} - price_{zjt}) + \beta_{zt}x_j + \xi_{jt} + \epsilon_{zjt} \tag{10}$$

from purchasing SKU $j \in J_{zt}$, where $income_{zt} - price_{zjt}$ represent a consumer's total expenditure on all other products. Hence, α_{zt} represents the marginal benefit to the consumer of consuming an extra dollar of the numeraire good and is predicted to have a negative sign.

Also included in the utility function is x_j, a vector of SKU-specific characteristics that consists of 15 flavor dummies, 6 volume dummies, 3 manufacturer dummies, and a dummy indicating whether SKUs have been bundled together. These dummies were generated from the information listed in table 2. The coefficients on the dummies β_{zt} represent the marginal benefit to the consumer of consuming a particular flavor, size, or manufacturer. Because size is assumed to be vertically differentiated with larger sizes preferred to smaller ones, the coefficients on larger sizes should be greater than the coefficients on smaller ones.

Berry (1994) advocates including ξ_{jt} in the utility specification to capture SKU characteristics that are unobserved by the econometrician but might be correlated with either α_{zt} or β_{zt}. Because sports drinks are basically flavored sugar water, it seems reasonable to believe that the flavor, size, and manufacturer dummies should capture SKU characteristics that are constant across weeks and zones, leaving only unobserved SKU characteristics that vary by week, time, or both.

To see this more formally, suppose $\xi_{jt} = \xi_j + \xi_t + \Delta\xi_{jt}$. This specification decomposes unobserved SKU characteristics into a SKU-specific component (ξ_j), a week-specific component (ξ_t) and a week-SKU interaction term $\Delta\xi_{jt}$. While the dummies described above capture ξ_j, they do not capture demand shocks embodied in ξ_t or promotional activities embodied in $\Delta\xi_{jt}$. Since both demand shocks and promotional activities are correlated with prices, without additional controls the price coefficient α_{zt} will be biased upwards.

Finally, the utility model includes ϵ_{zjt}, the random component which in a multinomial logit model is iid and follows an extreme value distribution. This distributional assumption is useful because when combined with the assumption that a shopper chooses the sports drink that maximizes utility, this distribution yields a closed form solution for the (conditional) probability that product j is chosen

$$Pr(j \text{ is chosen}|\gamma_{zt}, X) = \frac{\exp(-\alpha_{zt} price_{zjt} + \beta_{zt} x_j + \xi_{jt})}{\sum_{h=1}^{J_{zt}} \exp(-\alpha_{zt} price_{zht} + \beta_{zt} x_h + \xi_{ht})}. \tag{11}$$

In the above equation, $\gamma_{zt} = (\alpha_{zt}, \beta_{zt})$ are random coefficients with

$$\gamma_{zt} = \Delta' r_{zt} + \nu_{zt}, \nu_{st} \sim \text{iid } N(0, V_\gamma). \tag{12}$$

Equation 12 asserts that the vector γ_{zt} is normally distributed with mean $\Delta' r_{zt}$ and a variance-covariance matrix V_γ. The vector r_{zt} contains zone-specific information including the number of customers who visit stores within a zone during a particular week as well as the total number of brands stocked by stores within a zone during that week. These two variables proxy for the unobserved demand shock embodied in ξ_t. Also included in r_{zt} are a series of zone dummies that capture time-invariant characteristics unique to a pricing zone.

4.2.2 Choosing the Outside Good

Because all choice probabilities must sum to one, only the coefficients on $j-1$ of the products are identified. The traditional solution to this problem is to treat one of the goods as the "outside option", or the product that wasn't chosen, and set all the coefficients for this good equal to zero. Doing this implies that coefficients on the $j-1$ other products are measured relative to the outside good.

Here, I face a number of options for the outside good. One option is to use the difference between the number of customers who visit a store in a week and the number of sports drink SKUs sold in that store during that week. Unfortunately, because the sports drink sales constitute such a small fraction of all products sold at a supermarket, the probability of choosing the outside good would be close to one, which leads to identification problems.

Another alternative would be to to use the market for spring water as the outside good. According to internal memos, prior to Coca-Cola's entry into the market, Gatorade considered bottled water to be its primary competitor, making it an excellent choice for an outside good. Unfortunately, Dominick's stores do not start selling bottled spring water until almost halfway through the sample[6].

Ultimately, I decided to use the 32 ounce lemon-lime Gatorade as the outside good. Since virtually every store sells this SKU each week, I don't run afoul of the problems plaguing spring water. The main disadvantage of using this SKU as the outside good is that all the coefficients in the utility model must be interpreted relative to an SKU within the market. As a result, the coefficients on the 32 ounce, lemon-lime, and Gatorade dummies are no longer identified and must be omitted from the utility model.

4.2.3 The Bayesian Method

I use the methods described in Rossi et al. to simulate draws from the joint distribution of $\gamma = (\gamma_{11}, \ldots, \gamma_{1t}, \ldots, \gamma_{s1}, \ldots, \gamma_{st})$ and Δ in equation 11. With these draws in hand, I use the

[6] Dominick's did sell tonic, seltzer, and sparkling water throughout the sample. Tonic and seltzer water are primary ingredients in alcoholic cocktails and so are not likely to be close substitutes for Gatorade. I was unsure whether sparkling water was a close substitute for Gatorade.

means of the marginal distributions for my analysis. These means are comparable to (but not necessarily equal to) the multinomial logit's maximum likelihood estimator. Indeed, the main difference between the two estimators is philosophical; the Bayesian method treat α, β and Δ in equation 11 as random variables whose joint distribution is to be constructed, while maximum likelihood estimators treat α β and Δ as unknown constants whose values can be inferred from the asymptotic distribution of the maximum likelihood estimator. Of course, this philosophical distinction lead to practical differences in implementing these two methods. The Bayesian method constructs the joint distribution of α β and Δ using Bayes' rule, which is essentially the product of the likelihood function and a prior on α, β and Δ. Maximum likelihood methods construct their best guess of α, β, and Δs by calculating the mode of the likelihood function. See appendix A for the full derivation of the Bayesian estimator used here.

The Bayesian method has some advantages over maximum likelihood methods. First, while both methods require the econometrician to specify a likelihood function, only maximum likelihood attempts to find the likelihood's mode. The numerical procedures for locating the mode can fail, particularly when the likelihood function has either many local maxima or flat regions. This problem manifests itself as the "sensitivity to initial conditions" issue, where, where the initial values used in the numerical maximization procedures often determine whether the likelihood's mode is actually reached. In contrast, the Bayesian method integrates over the distribution of α and β to form an estimator While this estimator may not be informative, it is certainly computationally feasible. Second, the maximum likelihood estimator for the multinomial logit does not necessarily exist when some shoppers don't purchase all of the alternatives that are available. Multiplying the the likelihood by even a vague prior, however, ensures that a the Bayesian method yields a solution.

5 Results

In section 2, I argued that the trade-off between discrimination and differentiation could be uncovered by observing how a product's introduction or retirement affected the price menu of an existing product. I demonstrated that if a product was a sufficiently close substitute, it's introduction (removal) would increase (decrease) the value of the outside option, altering the price schedule and changing the profits that a supermarket earns on an existing product.

Here, I present the results from the estimation strategies discussed earlier. First, I present the differences-in-differences results under the assumption that all product introduction and retirements affect existing products in a similar fashion. Next, I present the estimation results from the random coefficient multinomial logit model and use these results to show how an incumbent SKU's own demand elasticities change when a new product is introduced. Finally, I re-estimate the differences-in-differences, this time including the change in an existing SKU's own price elasticity to account for the possibility that the introduction of a new product affects some existing SKUs differently than others.

5.1 Differences-in-Differences

To estimate the effect of product entry (exit), I identify the week in which a product first (last) enters (exits) the Dominick's chain and then execute regression 9 on stores that adopt (retire) the product within 10 weeks of that initial (terminal) date. Doing so excludes stores who are either late adopters (retirees) or never adopt (retire) the product.

I plot the coefficients on the relative time dummies (δ) against weeks to release (retirement) in figure 4. The solid blue line plots the regression coefficients, while the dashed brown lines surrounding the blue display the 95% confidence interval. The gap at week -1 is meant to ease interpretation of the graphs; omitting δ_{-1} implies that the remaining time dummies are interpreted relative to week -1, making δ_0 the estimate the causal effect of entry (exit) on retail price.

Figure 4(a) displays the effect of a new product entry on an existing product[7]. Four weeks before a new product enters a store, relative markups grew at 1%. Entry causes a 1% decrease in markups when the product first enters a store, followed by a 2-4% decrease in markups in subsequent weeks.

Similar effects hold for product exit. Prior to the product's exit, the markups on all other products decline by 1%. per week. After the product is removed, however, markups do not change. Unfortunately, the results are not significant at the 5% level.

Taken together, these results indicate that any trade-off between price discrimination and prod-

[7]These graphs only display δ. The full regression results are available from the author upon request.

uct differentiation is at best modest. This result is consistent with the theoretical predictions described in section 2, provided that the products added or removed from a supermarket's shelves are poor substitutes for the remaining products. The theory predicts that the introduction or removal of a poor substitute will do little to change the value of the outside option and hence will have only a small effect on price.

It seems unlikely, however, that all of the products that were introduced or removed were poor substitutes for existing products. Indeed, of the 19 new products that entered Dominick's stores 7 differed from existing products only by flavor, while 9 others differed only by their manufacturer. As I discussed earlier, the inability of regression 9 to account for these substitution patterns could cause the DIFF strategy to understate the causal effect of entry (exit) on retail price.

5.2 The Bayesian Method

To better control for the different effects of entering (exiting) products, I estimate the random coefficients multinomial logit model described in section 4. The first column in Table 5 reports the mean and standard deviation[8] of γ, the random coefficients in the utility model described by equation 10, averaged across zones and weeks. The estimates in column 1 match some of the theoretical predictions discussed earlier; the coefficient on price is negative, accurately reflecting the dis-utility a consumer attains from exchanging income for a sports drink. Likewise, the volume coefficients reveal that consumers receive higher utility from larger sized products. Although theory makes no specific predictions regarding the coefficients on the flavor and manufacturer dummies, the coefficients nevertheless tell an interesting story. The fact that all of the flavor coefficients are negative suggests that Lemon-Lime is the flavor that yields the highest utility. New flavors, like Grape, Ice Tea, Mountain Blast, etc. tend to yield less utility than some of Gatorade's older flavors, like Orange and Fruit Punch. The coefficients on the manufacturer dummies reveal that Gatorade is the preferred flavor, followed by All-Sport and Powerade.

The second column in table 5 indicates that consumers enjoy variety; as the number of SKUs carried by a particular zone increases, the disutility that the typical consumer obtains from pur-

[8]Care should be taken in interpreting the standard deviations. the posterior distributions of γ and δ are not normally distributed, which means that the standard t-test cannot be applied.

chasing an SKU decreases. Adding new products has mixed effects on the utility that a consumer obtains from a particular flavor. Apparently, adding new SKUs makes the Citrus Cool, Mountain Blast, Orange, Tidal Burst, and Tropical Fruit flavors more attractive to the typical consumer than the Lemon-Lime flavor, while many of the remaining flavors become less attractive.

The addition of new SKUs also effects how the typical consumer values different sizes of a particular product. As the number of SKUs grows, the typical consumer values the 20, 32, and 46 ounce containers more than the 16, 64, and 128 ounce container sizes. This finding suggests that consumer preferences over size become more homogeneous as the number of SKUs present in a particular zone increase.

The third column in table 5 suggests that positive demand shocks like zone-level sales or warm weather raise the dis-utility that a consumer receives from purchasing a sports drink. Positive demand shocks also tend to make the typical consumer value the lemon-lime flavor over the other flavors and the 32 ounce size over the other sizes. These shocks also effect a consumer's brand preference; All-Sport becomes relatively more attractive than Gatorade, while Powerade becomes less attractive.

The remaining columns capture how consumer tastes vary by zone. The first row suggests that only consumers from zones 4, 12 and 16 obtain substantially more disutility from purchasing a sports drink SKU than the omitted zone. The remaining rows suggest that consumers in different zones do value flavors differently from one another. For instance, while consumers in zones 6, 7, 11, and 14 typically value the lemon-ice flavor less than consumers in the omitted zone, the coefficient for zone 4 customers is double that of zone 5 customers and 10 times that of zone 8 customers.

5.2.1 Demand Elasticities

Using the estimates that I obtain above, I calculate the own price elasticities for each SKU sold in a zone for the week before and after a new product is introduced into that zone and average the change in own-price elasticities across zones and SKUs that are part of the same product. Table 6 displays the average percentage change in the own price elasticities for all products in the subsample (rows), for each new flavor introduction (columns), with the median changes reported in the last row and column of table 6. These results indicate that own price elasticities are substantially affected by the introduction of new flavors; as the last row indicates, most flavor releases are associated with an elasticity increase from 12.3 to 273 percent.

The only exceptions to this trend are the Tropical Fruit and Watermelon releases, where most product elasticities decrease. Since first differencing does not control for other factors that may have changed demand elasticities[9], the elasticity decrease could be due to demand changes not associated with the product release.

5.3 Revisiting Differences-in-Differences

I re-estimate 9, this time including the demand elasticity estimates described above. The results, which are reported in figure 5, reveal that the markup estimates do not significantly change after controlling for changes in the demand elasticities. This suggests that store-specific week effects that are not present in the data are also not correlated with the dependent variable.

6 Conclusion

While the introduction of new products can increase profits by attracting new consumers, it can also weaken a retailer's ability to price discriminate. Using supermarket-level data on sports drinks, I first establish that stores engage in horizontal differentiation as well as price discrimination. I then investigate the effect of product entry and exit on the price schedule of existing products to determine whether stores exchange some of their ability to price discriminate for the ability to differentiate. Sports drinks are well-suited for this study because i) their product characteristics are observed in the data and ii) a number of new sports drinks enter stores over the course of the sample.

I employ a difference-in-differences strategy to investigate this model. I discover evidence supporting the notion that firm's with market power exchange discrimination for differentiation. This evidence, however, suggests that this trade-off is at best a second order effect.

My results contrast sharply with the literature on the relationship between competition and price

[9]See section 4 for a more detailed discussion.

discrimination. One simple explanation for these differences is that the multi-product supermarkets in my study can choose which products to carry as well as how to price these products, mitigating the trade-off between horizontal differentiation and price discrimination. Retailers in these competition studies, however, must take their competitor's prices and product offerings into account, limiting their ability to control the trade-off.

Throughout this paper, I have focused on the effect of product entry and exit on the ability of a retailer to price discriminate. An equally interesting topic is the effect of price discrimination on the retailer's ability to horizontally differentiate. One way to capture this effect is to modify the model presented in section 2 to include a shelf-space constraint. A binding shelf-space constraint captures the opportunity cost to the supermarket of substituting a new size of an existing product (price discrimination) for an existing product (horizontal differentiation). Unfortunately, the Dominick's data does not contain any information on the amount of shelf space in a particular store, which means that either additional data or stronger assumptions are needed to incorporate the shelf-space constraint into an estimation strategy.

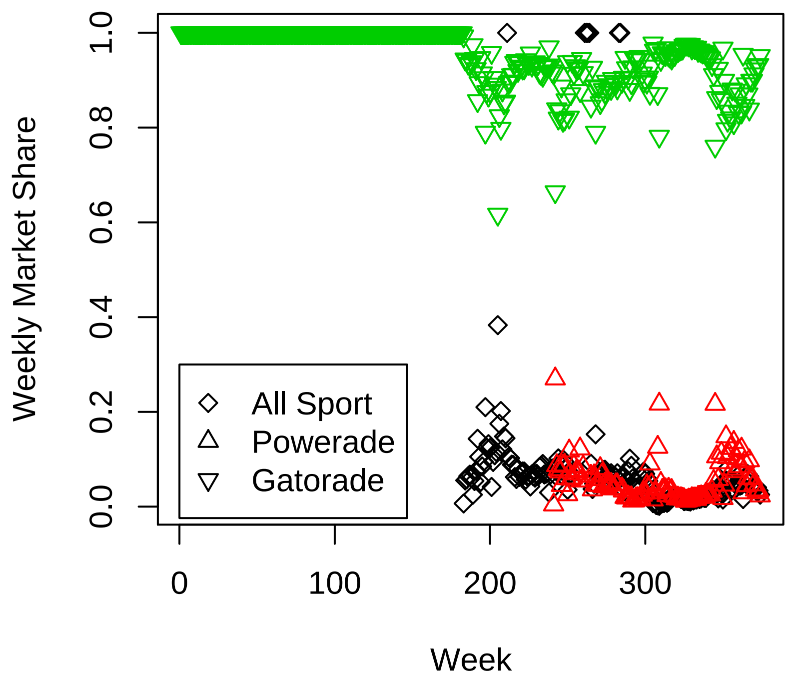

Figure 1: Sample Market Shares, 1989-1996

Year	Total Sales (Millions $)	All-Sport	Gatorade	Powerade
1988	474		100%	
1989	568		100%	
1990	676		100%	
1991	800		100%	
1992	800	Not Aval.		
1993	875	2.9%	82.4%	5.9%
1994	1000	7.6%	73.8%	10.5%
1995	1240	9.8%	72.3%	12.1
1996	1390	10.2%	72%	12.9%
1997	1480	9.6%	73.1%	14.3%

Table 1: U.S Sports Beverage Market, 1988-1997

Manufacturer	Flavor	Volume	Units Bundled	First Week	Final Week
Gatorade	Fruit Punch	16	4	1	41
Gatorade	Fruit Punch	20	1	287	348
Gatorade	Fruit Punch	20	4	337	375
Gatorade	Fruit Punch	32	1	1	375
Gatorade	Fruit Punch	64	1	1	375
Gatorade	Lemonade	16	4	1	53
Gatorade	Lemonade	46	1	1	240
Gatorade	Lemonade	64	1	1	375
Gatorade	Lemon-Lime	16	4	1	196
Gatorade	Lemon-Lime	20	4	335	375
Gatorade	Lemon-Lime	32	1	1	375
Gatorade	Lemon-Lime	46	1	1	292
Gatorade	Lemon-Lime	64	1	1	375
Gatorade	Lemon-Lime	128	1	1	78
Gatorade	Lemon-Lime	128	1	39	375
Gatorade	Orange	16	4	1	39
Gatorade	Orange	32	1	1	375
Gatorade	Orange	64	1	1	375
Gatorade	Orange	128	1	92	375
Gatorade	Citrus Cool	32	1	1	270
Gatorade	Citrus Cool	46	1	1	221
Gatorade	Citrus Cool	64	1	186	375
Gatorade	Lemon Ice	20	4	335	375
Gatorade	Lemon Ice	32	1	1	375
Gatorade	Lemon Ice	128	1	288	375
Gatorade	Tropical Fruit	32	1	82	374
Gatorade	Tropical Fruit	46	1	82	225
Gatorade	Tropical Fruit	64	1	133	375
Gatorade	Tropical Fruit	128	1	294	375
Gatorade	Grape	32	1	126	371
Gatorade	Grape	64	1	132	375
Gatorade	Grape	128	1	290	348
Gatorade	Iced Tea	32	1	186	312
Gatorade	Iced Tea	64	1	184	307
Gatorade	Watermelon	32	1	287	375
Gatorade	Watermelon	64	1	287	375
Gatorade	Raspberry	20	1	287	347
Gatorade	Raspberry	20	4	335	375
Gatorade	Raspberry	32	1	308	375
Gatorade	Raspberry	64	1	336	375
Gatorade	Strawberry Kiwi	20	4	335	375
Gatorade	Cherry	20	4	344	375
All-Sport	Fruit Punch	16	1	290	336
All-Sport	Fruit Punch	16	1	183	297
All-Sport	Fruit Punch	20	1	317	375
All-Sport	Fruit Punch	32	1	183	297
All-Sport	Lemon-Lime	16	1	183	297
All-Sport	Lemon-Lime	20	1	317	375
All-Sport	Lemon-Lime	32	1	183	297
All-Sport	Orange	16	1	183	297
All-Sport	Orange	16	1	290	315
All-Sport	Orange	20	1	309	375
All-Sport	Orange	32	1	183	297
All-Sport	Grape	20	1	309	375
All-Sport	Raspberry	20	1	342	375
All-Sport	Raspberry	32	1	343	375
All-Sport	Cherry	20	1	342	375
All-Sport	Cherry	32	1	343	375
Powerade	Fruit Punch	32	1	242	375
Powerade	Fruit Punch	64	1	249	300
Powerade	Grape	32	1	242	300
Powerade	Grape	64	1	250	300
Powerade	Lemon-Lime	32	1	242	375
Powerade	Lemon-Lime	64	1	249	300
Powerade	Orange	32	1	242	357
Powerade	Mountain Blast	32	1	310	375
Powerade	Tidal Burst	32	1	344	375

Table 2: Sports drink SKUs sold by Dominicks

Manufacturer	Flavor	Adopters	1	2	3	4	5	6	7	8	9	10
Gatorade	Tropical Fruit	80	0.29	0.82	0.94	0.97	1.00	1.00	1.00	1.00	1.00	1.00
Gatorade	Grape	84	0.31	0.98	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Gatorade	Ice Tea	85	0.01	0.79	0.96	0.98	1.00	1.00	1.00	1.00	1.00	1.00
Gatorade	Watermelon	78	0.28	0.65	0.86	0.88	0.91	0.99	0.99	0.99	1.00	1.00
Gatorade	Raspberry	68	0.81	0.91	0.97	0.99	0.99	0.99	0.99	1.00	1.00	1.00
Gatorade	Strawberry Kiwi	51	0.61	0.86	0.88	0.94	0.94	0.96	0.98	0.98	1.00	1.00
Gatorade	Cherry	61	0.56	0.72	0.85	0.90	0.93	0.95	0.95	0.95	0.95	0.95
All-Sport	Fruit Punch	73	0.48	0.58	0.68	0.77	0.79	0.82	0.86	0.90	0.97	0.99
All-Sport	Lemon-Lime	73	0.49	0.56	0.67	0.75	0.79	0.82	0.86	0.89	0.97	0.99
All-Sport	Orange	75	0.49	0.60	0.69	0.79	0.83	0.84	0.85	0.91	0.97	0.99
All-Sport	Grape	53	0.09	0.34	0.42	0.51	0.60	0.66	0.72	0.75	0.77	0.83
All-Sport	Raspberry	70	0.30	0.56	0.63	0.67	0.69	0.71	0.74	0.74	0.81	0.84
All-Sport	Cherry	71	0.17	0.37	0.48	0.58	0.65	0.68	0.70	0.75	0.79	0.83
Powerade	Fruit Punch	76	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Powerade	Grape	81	0.90	0.91	0.91	0.93	0.93	0.93	0.96	0.96	0.98	0.99
Powerade	Lemon-Lime	79	0.97	0.97	0.99	0.99	0.99	0.99	0.99	0.99	0.99	1.00
Powerade	Orange	75	0.96	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Powerade	Mountain Blast	68	0.09	0.12	0.13	0.25	0.41	0.47	0.50	0.60	0.60	0.66
Powerade	Tidal Burst	75	0.08	0.24	0.28	0.31	0.32	0.55	0.56	0.59	0.67	0.88
			0.47	**0.68**	**0.75**	**0.80**	**0.83**	**0.86**	**0.88**	**0.89**	**0.92**	**0.94**

Table 3: Fraction of Adopters in the Weeks Following a Product Introduction

Manufacturer	Flavor	Retirees	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1
Gatorade	Fruit Punch	15							0.67	0.07	0.07	0.07
Gatorade	Lemonade	44				0.84	0.84	0.82	0.80	0.80	0.07	0.05
Gatorade	Lemon-Lime	16							0.69	0.19	0.19	0.12
Gatorade	Orange	19							0.68	0.68	0.26	0.11
Gatorade	Citrus Cool	57						0.89	0.86	0.86	0.05	0.02
Gatorade	Lemon Ice	31							0.81	0.42	0.35	0.32
Gatorade	Tropical Fruit	46							0.87	0.30	0.26	0.20
Gatorade	Grape	52		0.88	0.88	0.88	0.88	0.87	0.79	0.15	0.12	0.08
Gatorade	Ice Tea	81					0.05	0.04	0.04	0.04	0.04	0.04
Gatorade	Watermelon	34							1.00	1.00	0.09	0.03
Gatorade	Raspberry	28						1.00	0.96	0.36	0.32	0.25
Gatorade	Strawberry Kiwi	62		0.97	0.95	0.94	0.94	0.92	0.82	0.31	0.29	0.18
Gatorade	Cherry	59				0.88	0.88	0.86	0.78	0.10	0.08	0.03
All-Sport	Fruit Punch	69			0.57	0.57	0.57	0.55	0.52	0.25	0.14	0.10
All-Sport	Lemon-Lime	66	0.56	0.56	0.56	0.55	0.55	0.52	0.45	0.12	0.09	0.06
All-Sport	Orange	72			0.49	0.47	0.47	0.46	0.39	0.19	0.14	0.08
All-Sport	Grape	48	0.67	0.67	0.67	0.62	0.62	0.56	0.52	0.38	0.31	0.15
All-Sport	Raspberry	51			1.00	1.00	1.00	0.98	0.92	0.41	0.37	0.20
All-Sport	Cherry	55		0.96	0.96	0.96	0.96	0.95	0.82	0.44	0.31	0.22
Powerade	Fruit Punch	68				0.94	0.94	0.91	0.85	0.19	0.16	0.15
Powerade	Grape	73	0.88	0.86	0.84	0.82	0.82	0.81	0.74	0.67	0.55	0.40
Powerade	Lemon-Lime	61				0.95	0.95	0.93	0.90	0.30	0.28	0.18
Powerade	Orange	80			0.17	0.16	0.07	0.07	0.07	0.06	0.05	0.01
Powerade	Mountain Blast	63				0.95	0.95	0.94	0.86	0.35	0.32	0.24
Powerade	Tidal Burst	59						0.98	0.93	0.29	0.27	0.19
			0.70	**0.82**	**0.71**	**0.77**	**0.72**	**0.74**	**0.71**	**0.36**	**0.21**	**0.14**

Table 4: Fraction of Retirees in the Weeks Preceding a Product Removal

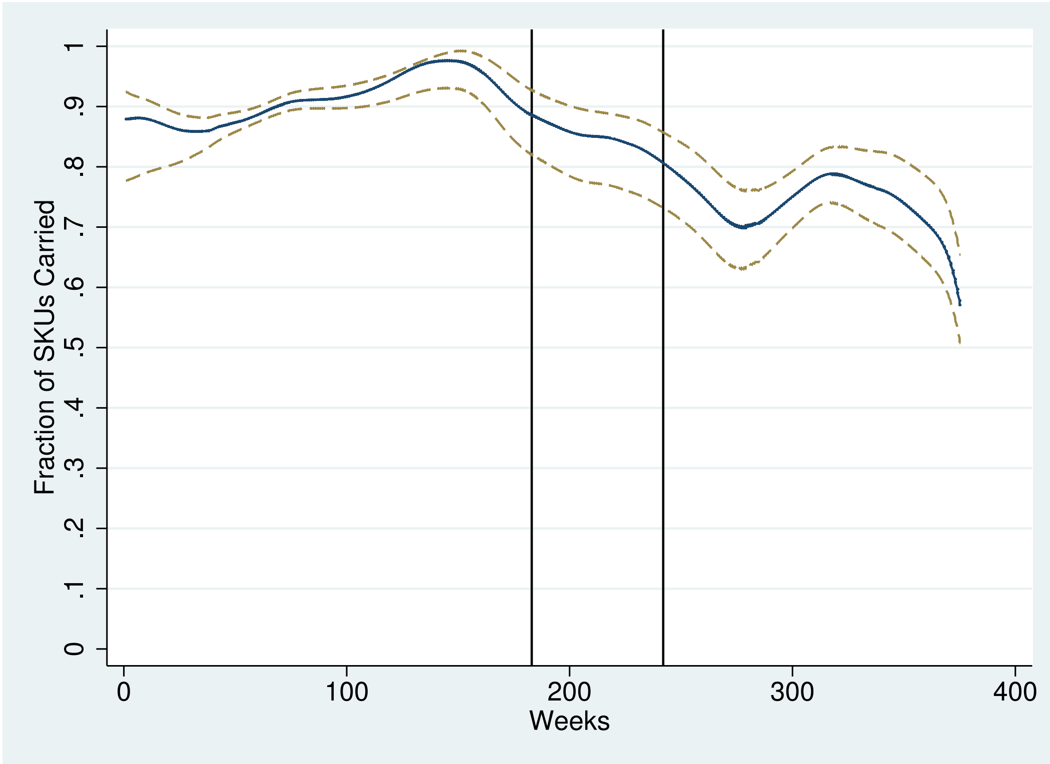

(a) SKUs

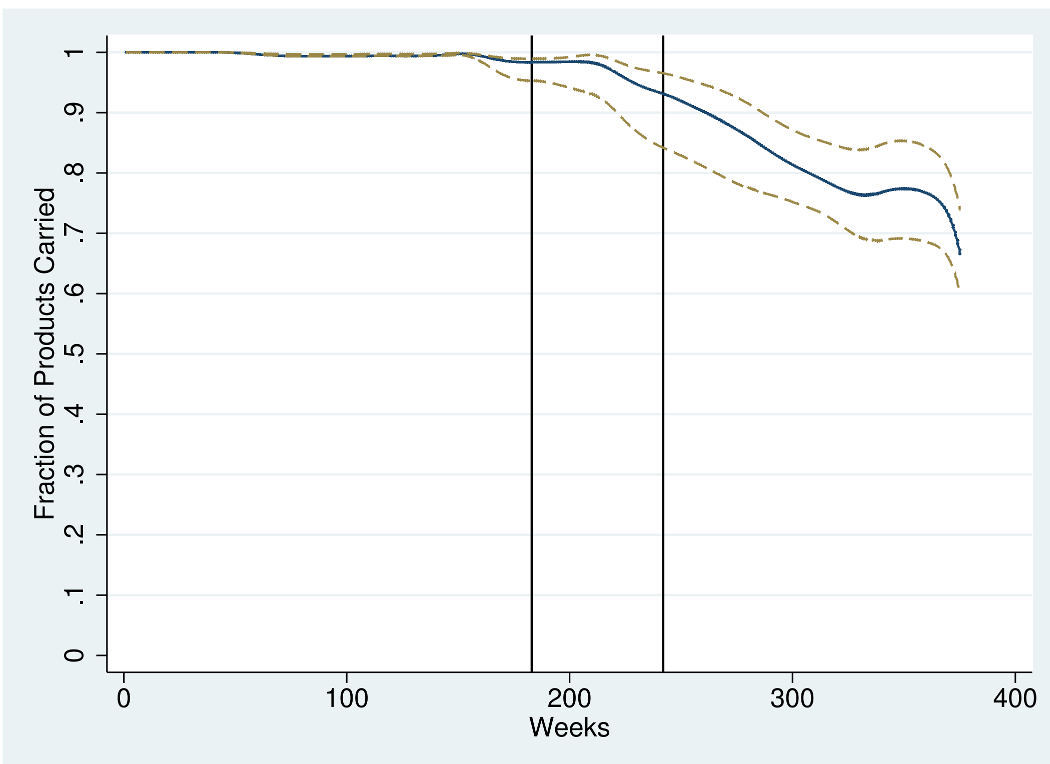

(b) Products

Figure 2: Fraction of SKUs and Products Carried by Stores Over Time

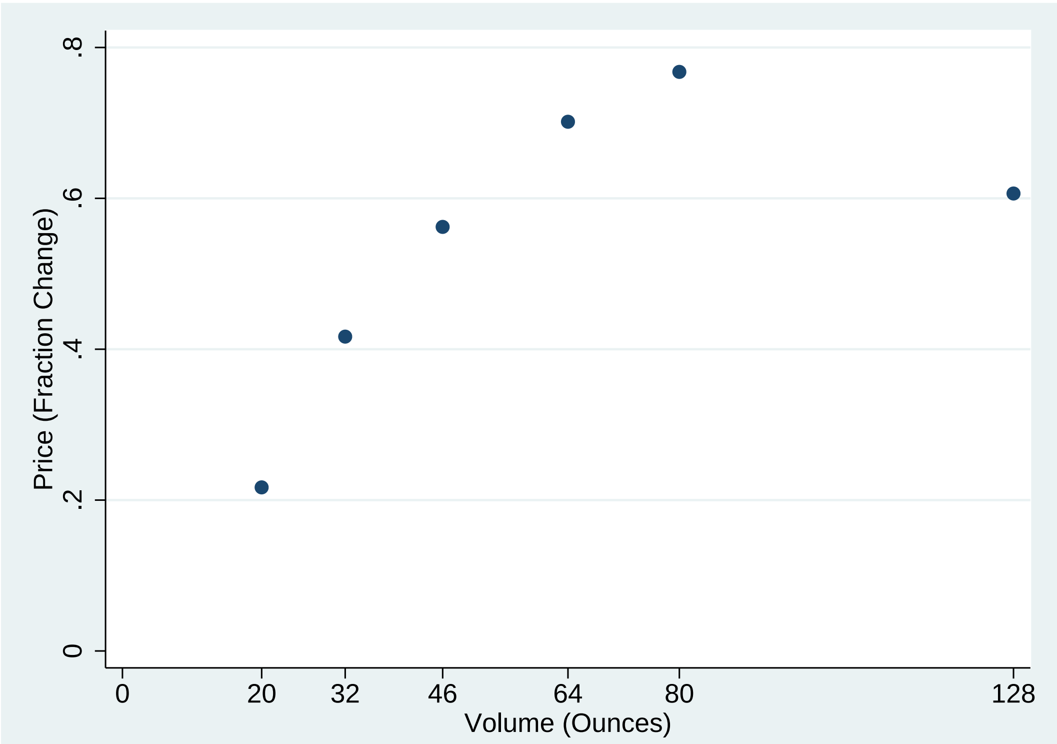

Figure 3: Price Discrimination in the Sports Drink Market

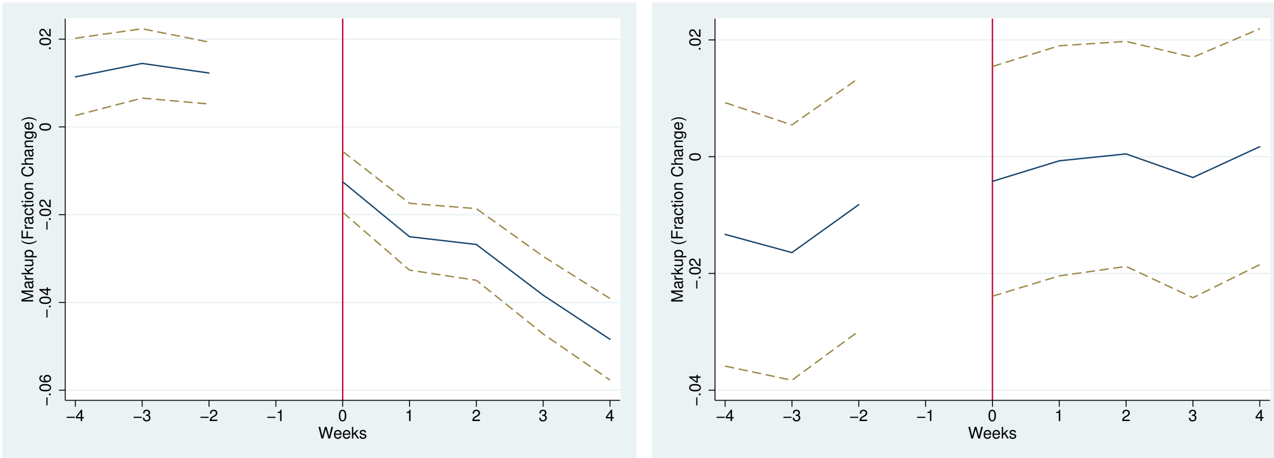

(a) New Product Enters Store

(b) Product Leaves Store

Figure 4: Product Entry and Removal

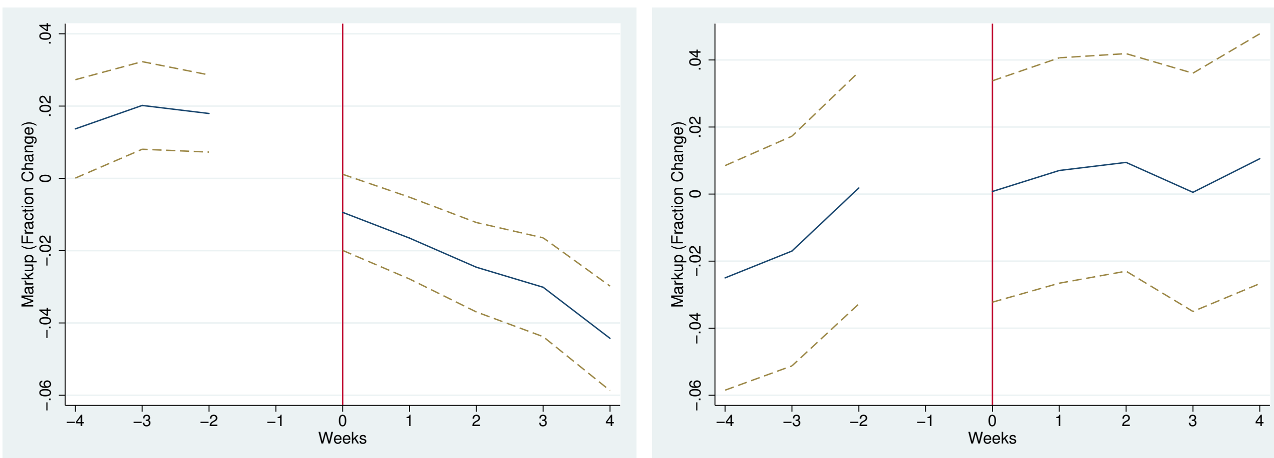

(a) New Product Enters Store

(b) Product Leaves Store

Figure 5: Product Entry and Removal, including Demand Elasticities

	γ	SKU Count	Customer Count	Zone 6	Zone 7	Zone 8	Zone 11	Zone 14	Zone 15
Price	-0.04	-0.0088	-0.011	-0.01	0.0036	0.0071	-0.0096	0.0008	-0.018
	(0.99)	(0.0082)	(0.022)	(0.0048)	(0.0034)	(0.0056)	(0.025)	(0.0045)	(0.0049)
Cherry	-0.1	-0.013	0.0058	0.0063	-0.0092	0.0047	0.0076	-0.012	-0.017
	(0.6)	(0.028)	(0.0034)	(0.013)	(0.032)	(0.0047)	(0.01)	(0.017)	(0.013)
Citrus Cool	-0.16	0.0067	0.01	0.041	-0.0023	-0.0025	$5.3e-08$	$-1.2e-06$	$-1.6e-07$
	(0.89)	(0.009)	(0.011)	(0.029)	(0.01)	(0.0097)	(2.4e-06)	(3.2e-06)	(2.2e-06)
Fruit Punch	-0.058	$5.2e-07$	$3.7e-06$	$6.3e-06$	$-1.3e-06$	$-1.2e-06$	$3.5e-06$	$1.1e-06$	$-3.2e-06$
	(0.6)	(1.4e-06)	(1.9e-06)	(7.2e-06)	(1.9e-06)	(2.2e-06)	(3.5e-06)	(1.4e-06)	(2.5e-06)
Grape	-0.11	$2.7e-06$	$-1.3e-07$	$-2.7e-06$	$-6.2e-06$	$-7.5e-07$	$-7.4e-07$	$-3.5e-06$	$-1.2e-06$
	(0.72)	(3.9e-06)	(1.8e-06)	(2.1e-06)	(4.3e-06)	(2.7e-06)	(3.1e-06)	(3.2e-06)	(8.2e-06)
Ice Tea	-0.24	$1.4e-06$	$-4.7e-06$	0.21	-0.19	0.023	0.031	0.69	0.77
	(0.69)	(2.1e-06)	(1.9e-06)	(0.35)	(0.51)	(0.34)	(0.22)	(0.3)	(1.2)
Lemon Ice	-0.16	-0.23	-0.22	0.9	0.21	-0.62	0.82	0.064	-0.47
	(0.82)	(0.3)	(0.34)	(0.55)	(0.23)	(0.4)	(0.62)	(0.29)	(0.33)
Lemonade	-0.24	-1	0.058	-0.25	-0.75	-0.87	0.24	-0.76	0.24
	(0.93)	(0.7)	(0.42)	(0.47)	(0.48)	(1.2)	(0.34)	(0.29)	(0.43)
Mountain Blast	-0.14	0.024	-0.011	0.02	0.85	1.1	-0.17	-0.24	1.1
	(0.8)	(0.59)	(0.4)	(0.26)	(0.36)	(1.4)	(0.35)	(0.4)	(0.62)
Orange	-0.043	0.24	-0.69	0.57	0.069	-0.37	-1.2	-0.061	-0.33
	(0.62)	(0.27)	(0.46)	(0.7)	(0.34)	(0.39)	(0.82)	(0.48)	(0.58)
Raspberry	-0.13	-0.91	-0.93	0.25	-0.9	0.19	-0.4	-0.014	0.049
	(0.85)	(0.58)	(1.5)	(0.4)	(0.34)	(0.59)	(0.84)	(0.56)	(0.36)
Tidal Burst	-0.13	1.1	1.9	-0.22	-0.27	1.1	0.27	-0.94	0.93
	(0.68)	(0.51)	(2.1)	(0.49)	(0.57)	(0.9)	(0.37)	(0.65)	(0.98)
Tropical Fruit	-0.1	0.031	-0.68	-1.7	0.043	-0.27	-1	-0.86	0.51
	(0.74)	(0.48)	(0.55)	(1.2)	(0.69)	(0.79)	(0.8)	(2)	(0.56)
Watermelon	-0.13	-1.1	0.4	-0.27	-0.11	0.085	1.1	1.6	-0.12
	(0.66)	(0.49)	(0.53)	(0.78)	(0.52)	(0.33)	(0.47)	(1.8)	(0.46)
16	-0.054	-0.36	1.4	0.36	-1.1	1.1	0.11	-0.6	-1.4
	(0.77)	(0.52)	(0.81)	(0.35)	(0.61)	(0.91)	(0.44)	(0.5)	(1.1)
20	-0.0091	0.048	-0.55	-1.3	-1.6	0.37	-1.1	0.081	-0.57
	(0.86)	(0.63)	(0.73)	(0.73)	(1.8)	(0.52)	(0.45)	(0.62)	(0.83)
46	0.0049	0.072	0.14	1.2	1.9	-0.11	-0.18	1.4	0.38
	(0.82)	(0.58)	(0.37)	(0.51)	(2)	(0.5)	(0.57)	(0.9)	(0.38)
64	0.029	-0.74	0.51	0.23	-0.65	-1.5	-0.033	-0.22	-0.98
	(1.4)	(0.66)	(1)	(0.49)	(0.55)	(1.2)	(0.71)	(0.84)	(0.84)
128	0.071	-0.51	0.37	-1.2	0.23	-0.29	0.14	0.11	1
	(3.5)	(2.1)	(0.56)	(0.49)	(0.5)	(0.71)	(0.48)	(0.31)	(0.43)
All-Sport	-0.048	1.9	-0.047	-0.21	1.4	0.33	-0.73	1	0.081
	(0.65)	(1.7)	(0.42)	(0.49)	(0.77)	(0.32)	(0.56)	(0.85)	(0.41)
Powerade	-0.12	-0.57	-1.2	0.064	-0.38	-0.92	-0.86	0.37	-1
	(0.6)	(0.47)	(0.99)	(0.59)	(0.68)	(0.68)	(1.7)	(0.48)	(0.42)

Table 5: Posterior Mean of γ, Δ

Manufacturer	Flavor	Tropical Fruit	Grape	Ice Tea	Watermelon	Mountain Blast	Cherry	Tidal Burst	
All-Sport	Cherry							699	**699**
Gatorade	Citrus Cool	−9.41	−0.928	80.2	11.3	−31.4	109	222	**11.3**
All-Sport	Fruit Punch			43.8	−35.8	540	104	287	**104**
Powerade	Fruit Punch				−13.4	81.6	131	1.03	**41.3**
Gatorade	Fruit Punch	−8.19	7.94	212	16	36.6	218	262	**36.6**
All-Sport	Grape					31.5	288	143	**143**
Powerade	Grape				10.7				**10.7**
Gatorade	Grape			114	−36.5	−3.42	3740	176	**114**
Gatorade	Ice Tea				2.55				**2.55**
Gatorade	Lemon Ice	−6.49	57.6	118	−51.7	118	171	219	**118**
All-Sport	Lemon-Lime			115	22.8		433	106	**110**
Powerade	Lemon-Lime				−16.3	14.3	33.1	971	**23.7**
Gatorade	Lemon-Lime	−22	12.3	67.1	21.1	102	36.1	409	**36.1**
Gatorade	Lemonade	−39.7	20.4	79.4	−52.8	13.8	208	375	**20.4**
Powerade	Mountain Blast						11.8	374	**193**
All-Sport	Orange			324	11.7	264	209	475	**264**
Powerade	Orange				4.17	558	1330	−37.3	**281**
Gatorade	Orange	−4.62	7.67	67.3	−29.9	−30.5	342	505	**7.67**
All-Sport	Raspberry							181	**181**
Gatorade	Raspberry					129	331	284	**284**
Gatorade	Tropical Fruit		21.8	45.9	−4.52	181	150	84.2	**65.1**
Gatorade	Watermelon					82.4	122	395	**122**
		−8.8	**12.3**	**80.2**	**−0.982**	**82**	**190**	**273**	

Table 6: Percentage Change in Own Price Elasticities Due to New flavor Entry

A The Bayesian Method

Here, I describe how I structurally estimate demand using Bayesian econometric methods.

A.1 Deriving the posterior

To implement the Bayesian method, I use Bayes' rule to express the conditional distribution of γ and Δ as a function of the likelihood and prior

$$\begin{aligned}
Pr(\gamma, \Delta | X, y, V_\gamma) &= \frac{Pr(X, y|\gamma, \Delta, V_\gamma)Pr(\gamma, \Delta, V_\gamma)}{Pr(X, y, V_\gamma)} \\
&\propto Pr(X, y|\gamma, \Delta)Pr(\gamma, \Delta, V_\gamma) \\
&= Pr(X|\phi)Pr(\phi)Pr(y|X, \gamma, \Delta, V_\gamma)Pr(\gamma, \Delta, V_\gamma) \qquad (13) \\
&\propto Pr(y|X, \gamma, \Delta, V_\gamma)Pr(\gamma, \Delta, V_\gamma) \\
&= Pr(y|X, \gamma)Pr(\gamma|\Delta)Pr(\Delta|V_\gamma)Pr(V_\gamma)
\end{aligned}$$

The first line in equation 13 is a direct application of Bayes' rule. The second line follows since the denominator is not a function of γ or Δ and can be recovered by integrating the numerator. The third line decomposes the joint probability of the independent and dependent variables X, y into an equivalent conditional and marginal probability, while the fourth line asserts that the term $Pr(X|\phi)Pr(\phi)$ is irrelevant if ϕ is independent of γ and Δ. The final line again applies Bayes' rule, this time to decompose $Pr(\gamma, \Delta, V_\gamma)$ into a series of conditional distributions. Notice how in this last equation, Δ, V_γ have disappeared from the first term, and the second term is not conditioned on V_γ; these terms are redundant given the other conditioning variables.

Equation 13 indicates that the joint distribution of γ, Δ is proportional to two terms; $Pr(y|X, \gamma)$ or the likelihood of observing the data conditional on the data, γ, and $Pr(\gamma, \Delta, V_\gamma)$, the econometrician's prior beliefs regarding the distribution of the random coefficients γ. The final line in equation 13 is important because it allows me to explicitly leverage the relationship between γ and Δ that is indicated by equation 12. Bayesians refer to the final line in equation 12 as a hierarchical model, with $Pr(\gamma|\Delta)$ as the first stage prior and $Pr(\Delta, |V_\gamma)Pr(V_\gamma)$ as the second stage prior.

Multiplying the likelihood by the prior effectively allows the econometrician's current beliefs regarding the distribution of γ, Δ to be updated by evidence. It is important to note however, that unless the econometrician adopts a fairly dogmatic prior (like restricting the joint density of γ, Δ to 0), that in many cases the evidence contained in the likelihood will overwhelm the prior. Because I do not have strong opinions regarding the distribution of γ, Δ, I use a fairly diffuse prior, which should be easily overwhelmed by the number of observations in the likelihood.

If I observed an individual shopper's purchases, the likelihood under the multinomial logit model for the N_{st} consumers who purchase sports drinks at store s during week t would be

$$\ell_{st}(\gamma_{st}) = \prod_{i=1}^{N_{st}} \prod_{j=1}^{J_{st}} Pr(j \text{ is chosen}|\gamma_{st}, X)^{1(j \text{ is chosen})} \tag{14}$$

where $1(j \text{ is chosen})$ is an indicator variable equal to 1 if the ith shopper chooses product j and 0 otherwise, and $Pr(j \text{ is chosen}|\gamma_{st}, X)$ is given by equation 11. Furthermore, because equation 12 implies that conditional on Δ , all the elements of γ are independent, I can rewrite the joint distribution of γ, Δ as

$$Pr(\gamma, \Delta|X, y, V_\gamma) = \prod_{s=1}^{S} \prod_{t=1}^{T} \ell_{st}(\gamma_{st}) Pr(\gamma_{st}|\Delta) Pr(\Delta|V_\gamma) Pr(V_\gamma) \tag{15}$$

Unfortunately, I do not observe individual shoppers. Instead, I observe the price and total sales for each product sold at every Dominick's supermarket. As such, conditional on a draw from γ_{st}'s distribution, shoppers at store s and week t are identical to one another, and the likelihood function for the N_{st} consumers who shop at store s during week t becomes

$$l_{st}(\gamma_{st}) = \prod_{j=1}^{J_{st}} n_{jst} Pr(j \text{ is chosen}|\gamma_{st}, X) \tag{16}$$

where n_{jst} is the number of shoppers who purchased product j in store s during week t and $\sum n_{jst} = N_{st}$. This yields a joint distribution for γ, Δ equalling

$$Pr(\gamma, \Delta|X, y, V_\gamma) = \prod_{s=1}^{S} \prod_{t=1}^{T} l_{st}(\gamma_{st}) Pr(\gamma_{st}|\Delta) Pr(\Delta|V_\gamma) Pr(V_\gamma) \tag{17}$$

All I have left to do is specify the densities for $\gamma_{st}|\Delta)$, $\Delta|V_\gamma$ and V_γ. In fact, $\gamma_{st}|\Delta$ has already been specified; equation 12 indicates that $\gamma_{st}|\Delta \sim N(\Delta' z_{st}, V_\gamma)$. To make sampling from the posterior easier, I will assume that $\text{vec}(\Delta)|V_\gamma \sim N(\text{vec}(\bar{\Delta}), V_\gamma \otimes A^{-1})$ and V_γ follows an inverse Wishart distribution with parameters ν, V.

A.2 Sampling from the posterior

To obtain the random coefficient multinomial logit estimator, I need to sample from the posterior described in equation 17. Unfortunately, the posterior distribution is complicated enough so that standard sampling methods available in most statistical software packages won't suffice. Instead, Rossi et al. describe a two step algorithm that I use to draw from the posterior distribution. This algorithm partitions the parameters into two groups, one containing Δ and V_γ and the other containing γ. The algorithm then describes how to sample from $\Delta, V_\gamma|\gamma$ and then $\gamma|\Delta, V_\gamma$. According to equation 13, this algorithm will yield a draw from the posterior distribution.

To construct Δ, and V_γ needed for the first step, Rossi et al. note that for a particular realization of γ, equation 12 describes a (Bayesian) multivariate regression model. Using the fact that

$$Pr(\Delta, V_\gamma|\gamma, r) = Pr(\Delta|V_\gamma, \gamma, r)Pr(v_\gamma|r)$$

$$r = (r_{11}, \ldots, r_{1t}, \ldots, r_{s1}, \ldots, r_{st})$$

Rossi et al. demonstrate that under the priors stated earlier,

$$\begin{aligned}
V_\gamma|\gamma, r &\sim IW(\nu + n, V + S), \\
\text{vec}(\Delta)|V_\gamma, \gamma, r &\sim N(\text{vec}(\tilde{\Delta}), V_\gamma \otimes (r'r + A)^{-1}), \\
&\quad \tilde{\Delta} = (r'r + A)^{-1}(W'\gamma + A\bar{\Delta}), \\
&\quad S = (\gamma - X\tilde{\Delta})'(\gamma - X\tilde{\Delta}) + (\tilde{\Delta} - \bar{\Delta})'A(\tilde{\Delta} - \bar{\Delta})
\end{aligned}$$

In other words, to obtain a draw from $\Delta, V_\gamma|\gamma, r$ first sample V_γ using the inverse Wishart distribution and then draw $\Delta|V_\gamma$ using the Normal distribution.

The second step in Rossi et al.'s algorithm is to sample from $\gamma|\Delta, V_\gamma$, or equivalently, sample from the likelihood described in equation 16. Rossi et al. recommend using a Gaussian Random-Walk Metropolis Markov Chain to obtain draws from the likelihood. Bayesians typically employ Metropolis Chains when, like in this case, the desired distribution cannot be sampled from using traditional sampling schemes. Like all Metropolis Chains, the Gaussian Random-Walk Metropolis chain starts with an initial value γ^0 and samples from a "candidate" distribution that is similar to the desired distribution but easier to sample from. Unlike other Metropolis chains, the Gaussian Random-Walk Metropolis chain assumes that the candidate draw θ is a Gaussian random walk

$$\theta = \gamma + \epsilon \quad , \epsilon \sim N(0, s^2\Sigma) \tag{18}$$

where s is a scaling factor and Σ is the variance-covariance matrix. Both the scaling factor and Σ must be chosen so that the Markov Chain's draws mimic the shape of the posterior distribution. Following Rossi et al., I choose $s = 2.93/\sqrt{(\dim(\gamma_{st})}$ and $\Sigma = (H_{st} + (V_\gamma^r)^{-1})^{-1}$, where H_{st} is the Hessian of the likelihood for store s and week t, evaluated at the maximum likelihood estimate that is obtained from the "fraction likelihood"[10], and V_γ^r is the current draw of V_γ from step 1. Rossi et al. recommend a Σ based on H_{st} because of the fact that the multinomial logit model is well approximated by a normal distribution with mean equal to the MLE of $l_{st}(\psi)$ and variance equal to the Hessian of the likelihood evaluated at the MLE of $l_{st}(\psi)$.

Once the candidate draw has been sampled, the Metropolis algorithm accepts the draw with probability $\alpha = \min\{1, \frac{\pi(\theta)}{\pi(\gamma)}\}$. Steps 1 and 2 are then repeated until the sampled draws behave as

[10]The fraction likelihood is $l_{st}(\psi)\bar{l}(\psi)^{N_{st}/cN}$, where $\bar{l}(\psi)^{N_{st}/cN}$ is the pooled likelihood scaled in proportion to the fraction of consumers at store s during week t. This procedure is necessary since not all sports drinks are sold at each stores, implying that the $l_{st}(\psi)$ might not have a maximum. Scaling $l_{st}(\psi)$ by $\bar{l}(\psi)^{N_{st}/cN}$ which does have a maximum, remedies the problem.

if they had been drawn from the posterior distribution.

References

Anderson, S. P., A. De Palma, and J. F. Thisse (1992). *Discrete choice theory of product differentiation*. Cambridge, Mass.: The MIT Press.

Anderson, S. P. and M. Engers (2001, March). Preemptive entry in differentiated product markets. *Economic Theory 17*(2), 419–445.

Armstrong, M. (1996, jan). Multiproduct nonlinear pricing. *Econometrica 64*(1), 51–75.

Berry, S., J. Levinsohn, and A. Pakes (1995, Jul.). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society 63*(4), 841–890.

Berry, S. T. (1994, Summer). Estimating discrete-choice models of product differentiation. *The Rand Journal of Economics 25*(2), 242–262.

Besanko, D., S. Gupta, and D. Jain (1998). Logit demand estimation under competitive pricing behavior: An equilibrium framework. *Management Science 44*(11), 1533–1547.

Borzekowski, R., R. Thomadsen, and C. Taragin (2006, July). Competition and price discrimination in the market for mailing lists.

Busse, M. and M. Rysman (2005, Summer). Competition and price discrimination in yellow pages advertising. *RAND Journal of Economics 36*(2), 378–390. available at http://ideas.repec.org/a/rje/randje/v36y20052p378-390.html.

Card, D. and A. B. Krueger (1994, sep). Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania. *The American Economic Review 84*(4), 772–793.

Chevalier, J. A., A. K. Kashyap, and P. E. Rossi (2000, October). Why don't prices rise during periods of peak demand? evidence from scanner data.

Chintagunta, P. K., J.-P. Dub, and V. Singh (2003). Balancing profitability and customer welfare in a supermarket chain. *Quantitative Marketing and Economics 1*(1), 111–147.

Cohen, A. (2004, January). Indentifying price discrimination when product menus are endogenous.

Draganska, M. and D. C. Jain (2005, March). Product-line length as a competitive tool. *Journal of Economics & Management Strategy 14*(1), 1.

Draganska, M. and D. C. Jain (2006). Consumer preferences and product-line pricing strategies: An empirical analysis. *MArketing Science 25*(2), 164–174.

Ellickson, P. B. and S. Misra (2006, May). Supermarket pricing strategies.

Goolsbee, A. and C. Syverson (2005, January). How do incumbents respond to the threat of entry? evidence from the major airlines.

Hoch, S. J., B.-D. Kim, A. L. Montgomery, and P. E. Rossi (1995, Feb.). Determinants of store-level price elasticity. *Journal of Marketing Research 32*(1), 17–29.

Hotelling, H. (1929, mar). Stability in competition. *The Economic Journal 39*(153), 41–57.

Itoh, M. (1983, October). Monopoly, product differentiation and economic welfare. *Journal of Economic Theory 31*(1), 88–104.

Jacobson, L. S., R. J. LaLonde, and D. G. Sullivan (1993, Sep.). Earnings losses of displaced workers. *The American Economic Review 83*(4), 685–709.

Judd, K. L. (1985, Summer). Credible spatial preemption. *The Rand Journal of Economics 16*(2), 153–166.

Kadiyali, V., N. Vilcassim, and P. Chintagunta (1999). Product line extensions and competitive market interactions: An empirical analysis. *Journal of Econometrics 89*(1-2), 339–363.

Kano, K. (2005, November). Modes of competition and asymmetric price rigidity in a menu-cost economy.

Lancaster, T. (2004). *An introduction to modern Bayesian econometrics*. Malden, MA: Blackwell Pub.

Levitt, S. D. (1997, jun). Using electoral cycles in police hiring to estimate the effect of police on crime. *The American Economic Review 87*(3), 270–290.

Maskin, E. and J. Riley (1984). Monopoly with incomplete information. *The RAND Journal of Economics 15*(2), 171–196.

Mcmanus, B. (2004, December). Nonlinear pricing in an oligopoly market: The case of specialty coffee.

Meza, S. and K. Sudhir (2005, December). The role of strategic pricing by retailers in the success of store brands.

Montgomery, A. L. (1997). Creating micro-marketing pricing strategies using supermarket scanner data. *Marketing Science 16*(4), 315–337.

Nevo, A. (2000, Winter). A practitioner's guide to estimation of random coefficients logit models of demand. *Journal of Economics & Management Strategy 9*(4), 513–548.

Nevo, A. (2001, Mar.). Measuring market power in the ready-to-eat cereal industry. *Econometrica 69*(2), 307–342.

Nevo, A. and K. Hatzitaskos (2005, July). Why does the average price of tuna fall during lent?

Prescott, E. C. and M. Visscher (1977, Autumn). Sequential location among firms with foresight. *The Bell Journal of Economics 8*(2), 378–393.

Rochet, J.-C. and L. A. Stole (2002, January). Nonlinear pricing with random participation. *The Review of Economic Studies 69*(1), 277–311.

Rossi, P. E., G. M. Allenby, and R. McCulloch (2005). *Bayesian Statistics and Marketing*. Wiley & Sons Inc.

Rovell, D. (2005). *First in Thirst: How Gatorade Turned the Science of Sweat into a Cultural Phenomenon*. American Management Association.

Salop, S. C. (1979, may). Strategic entry deterrence. *The American Economic Review 69*(2), 335–338.

Schmalensee, R. (1978). Entry deterrence in the ready-to-eat breakfast cereal industry. *Bell Journal of Economics 9*(2), 305–327. available at http://ideas.repec.org/a/rje/bellje/v9y1978iautumnp305-327.html.

Seim, K. and V. B. Viard (2004, November). The effect of entry and market structure on cellular pricing tactics.

Shepard, A. (1991, feb). Price discrimination and retail configuration. *The Journal of Political Economy 99*(1), 30–53.

Tirole, J. (1988). *The theory of industrial organization*. Cambridge, Mass.: MIT Press. Jean Tirole.; Includes bibliographies and index.

Train, K. (2003). *Discrete choice methods with simulation*. New York: Cambridge University Press.

Wooldridge, J. M. (2002). *Econometric analysis of cross section and panel data*. Cambridge, Mass.: MIT Press. Jeffrey M. Wooldridge.; Includes bibliographical references (p. [721]-735) and index.