Quantitative Marketing and Economics, 1, 111–147, 2003.

Balancing Profitability and Customer Welfare in a Supermarket Chain

PRADEEP K. CHINTAGUNTA*
University of Chicago, Graduate School of Business, 1101 E. 58th Street, Chicago, IL 60637, USA
Email: pradeep.chintagunta@gsb.uchicago.edu

JEAN-PIERRE DUBÉ
Graduate School of Business, University of Chicago, 1101 East 58 Street, Chicago, IL 60637
E-mail: jdube@gsb.uchicago.edu

VISHAL SINGH
Graduate School of Industrial Administration, Carnegie Mellon University, Schenley Park, Pittsburgh, PA 15213
E-mail: vsingh@andrew.cmu.edu

Abstract. We investigate the impact of price discrimination by a large Chicago supermarket chain. First we measure the impact of the chain's current zone-pricing policy on shelf prices, variable profits and consumer welfare across its stores. Using the chain's database to simulate a finer store-specific micro-pricing policy, we study the implications of this policy on profits and welfare. We show how a store-pricing policy that is constrained to offer consumers at least as much surplus as a uniform chain wide pricing policy still enables the retailer to generate substantial incremental profits.

To ensure our pricing problem exhibits a well-defined optimum, we use the parsimonious, mixed-logit demand function that allows for flexible substitution patterns across brands and also retains a link to consumer theory. We discuss the issue of price endogeneity when estimating the demand parameters with weekly store-level data. Standard instrumental variables techniques used to account for such endogeneity also seem to increase the magnitudes of own-price elasticities thereby offsetting the problem encountered by previous researchers of predicted prices from a demand model exceeding those in the actual data.

Key words. price discrimination, customer welfare, demand modeling

JEL Classification: D12, D4, L81, M3

1. Introduction

In recent years, the practice of "zone pricing" has become increasingly popular for retailers. Under this form of third-degree price discrimination a firm selects various delivered prices and the geographic zones in which they apply. In some instances, the definition of a zone may be sufficiently narrow that nearby outlets of a common

*Corresponding author.

retail chain may charge noticeably different prices. For instance, retail gasoline stations in a given city often charge higher prices at outlets near freeways and fast food chains charge different prices at airports. Access to a historic store-level database enables firms to consider even finer data-based micro-marketing policies that use the store-level information to recognize differences in underlying consumer demand. While researchers have documented the potential profit-gains from store-level pricing, the impact on consumer welfare has not been considered. A potential concern for chain managers interested in micro-marketing is the risk that appropriating too much surplus from consumers could, in the long-run, damage consumer relations or result in a loss of store traffic. Recognizing this customer relations problem, managers may be able to generate balanced price discrimination strategies across stores that generate additional profits without appropriating "too much" consumer surplus.

We investigate the role of a balanced micro-marketing pricing policy across stores from a large supermarket chain in the Chicago area—Dominick's Finer Foods (DFF). First, we estimate a system of demand equations within a product category. The derivation of the functional form for aggregate demand from consumer theory enables us to measure consumer welfare. We account for heterogeneity across consumers within a store by using a random coefficients specification. To capture heterogeneity across stores, we use store-level demographic information to characterize differences in the underlying consumer populations. In estimating the demand system, we also recognize how the limitations of our data, in the form of unmeasured "product characteristics," might generate an endogeneity problem. We apply an instrumental variables procedure that uses city-level wholesale prices as proxies for the shelf prices within a store. Combining the estimated demand system with a model of product category pricing for the retailer enables us to simulate various pricing policies. First, since we have information on the true margins, we are able to assess the validity of our model estimates in predicting actual observed prices at the chain. We then simulate prices under a variety of store pricing schemes and compute the resulting chain profits and consumer welfare.

Previous research discusses a problem encountered when estimating demand systems whereby the estimated own-price elasticities of demand generate optimal margin recommendations that exceed those observed in the data (Montgomery and Bradlow, 1999). Effectively, these elasticities are "too small." Several explanations have been suggested for this problem. Typically, this problem is attributed to the use of inappropriate functional forms of demand. For instance, double-log demand systems do not yield globally concave retail category profit functions, complicating the calculation of optimal prices (Reibstein and Gatignon, 1984; Anderson and Vilcassim, 2001). The use of a structural demand system ensures a well-behaved profit function so we do not need to impose constraints on the pricing problem (Montgomery, 1997) or to constrain the range of estimated demand parameters (Montgomery and Bradlow, 1999). Linear and double-log models also require a large number of parameters, which may lead to "noisy" or even incorrect-signed own and cross-price elasticities.[1] We use the parsimonious discrete choice demand system (Berry et al., 1995; Nevo, 2001). Allowing for heterogeneity across consumers within a store using a random coefficients specification generates flexible substitution patterns across brands with a relatively small number of parameters. Thus, our functional form is attractive since it constitutes a valid demand system and exhibits flexible substitution patterns using a relatively small number of parameters.

In addition to functional form problems, understated own-price elasticities using weekly supermarket data have also been attributed to price endogeneity due to unmeasured product characteristics (Besanko et al., 1998). We account for price endogeneity via an instrumental variables procedure. The instruments appear to lead to much more elastic estimates of demand. A final explanation for understated price elasticities may also be omitted sources of retail competition; although this seems less plausible given the results in Hoch et al. (1995). Moreover, Dreze et al. (1994b) run an experiment in the same DFF chain to show that prices could be raised permanently across stores without damaging store traffic.[2] As in Hoch et al. (1995), we include local competitive variables in our demand specification to account for retail competition. Our results indicate that the proposed discrete choice demand system with instrumented prices, and with local competitive variables predicts margins that are close to the levels observed in the data. This gives us some assurance that the demand model is capable of providing a realistic depiction of actual pricing decisions made by the chain.

Our empirical results indicate that both the price-sensitivity and the intrinsic preference to make a category purchase vary across stores. As in Hoch et al. (1995), we find that demographic variables consistent with willingness-to-pay are the most influential for market shares and elasticities. Thus, we conclude that the extant zone classification is based primarily on discriminating across consumer types, as opposed to competition or consumer search. Using the model estimates, we then compute the prices and quantities that would prevail if the stores adopted a store-specific pricing policy. Interestingly, while consumer welfare effects vary across stores, we note that store-pricing seems to target higher prices to less-affluent areas with larger ethnic populations and higher search costs. We also observe that store-level pricing may be better suited to some categories than others. For example, in the liquid laundry detergent category, while we find only modest gains from store level pricing, consumers are actually better off with such pricing. On the other hand, in the orange juice category, we do find sizeable losses in consumer welfare, if the chain moves to store pricing emphasizing the importance

1 Note that previous work has often estimated negative cross-price elasticities of demand in categories which one would expect to consist of substitute products. These incorrect signs will bias simulated profit-maximizing prices upwards.

2 They find no impact on store traffic after a three month period.

of balancing chain profits with consumer welfare when making micro-marketing pricing decisions.

Two practical considerations arise with such profit-enhancing pricing policies. First, a store manager could be concerned about the overall losses in consumer welfare in stores where prices are expected to rise. To mitigate these losses, we propose a constrained pricing procedure that makes use of the underlying economic structure of the model. The approach restricts the new store prices to offer the population of consumers in each store at least the same level of welfare as under a uniform chainwide pricing policy. We find that this constraint still enables the store manager to capture a large portion of the gains from unconstrained pricing in both categories. The second practical consideration involves the potential complexity of computing store-level prices for a chain which, like Dominicks, has a large number of stores. We address this computational complexity by suggesting an improved zone structure. Using our store-specific price levels, we cluster the stores into five zones. We find that these zones offer higher profits than the existing zone configuration.

1.1. Related literature

Two papers using the same data are closely-related to our work. Hoch et al. (1995) find that a large proportion of the variation in category-level consumer price elasticities across stores is explained by local consumer demographics and, to a much lesser extent, by local competitive variables. This result suggests that current zone-pricing by the chain appears to be driven more by a price discrimination motive than by competition. In a follow-up study, Montgomery (1997) looks at the profit-implications of zone and store pricing for the supermarket chain in one specific product category. This latter study documents the value to a retailer of implementing a data-based pricing policy. In particular, the differences in demand across stores within a given metropolitan area may be sufficient to enable profitable price discrimination. Our work differs from these two papers in several ways. Most importantly, we use a structurally-derived demand system. As a result, we are able to measure consumer surplus, an economic metric with which we can evaluate micro-marketing pricing strategies in conjunction with retail profits. The approach also has a number of practical advantages, such as the parsimony in terms of the number of parameters, existence of a well-defined optimal category price vector and the ability to measure consumer surplus.

Our work is also related to broader theoretical and empirical studies of price discrimination. First, our data permit us to rule out alternative explanations of the observed price dispersion across stores such as cost differences, consumer search and competition (e.g., Shepard, 1991; Borenstein, 1991; Cohen, 2001; Zhao, 2001). We then use a structural model to help us measure the micro-marketing prices and their corresponding profit and welfare implications (Leslie, 2001; Cohen, 2000). In the absence of competition, the theoretical profit implications of third-degree price discrimination are well-documented.[3] However, welfare gains are ambiguous; although a necessary condition for aggregate welfare gains in the single-product case requires aggregate unit sales to increase (Schmalensee, 1981; Varian, 1985). In our multiproduct category pricing context, the results are more ambiguous. In general, welfare gains may accrue, irrespective of category expansion, through the realignment of prices to match heterogeneity in consumer tastes across stores for the set of goods in the product line.[4] More importantly, consumer surplus may rise or fall (see Chiang and Spatt, 1982 for the case of a finite number of consumer types). Ultimately, predicting welfare gains or losses becomes an empirical issue which we address by explicitly measuring consumer surplus.

The rest of the paper is structured as follows. Section 2 presents the model. Section 3 discusses the main aspects of the estimation procedure. Section 4 provides an overview of the data we use for estimation. Section 5 presents results from the estimation of the demand parameters. Section 6 discusses the effects of micro-marketing pricing and the welfare implications of zone and store pricing. We conclude in Section 7.

2. Model

To develop a viable framework for a category pricing manager, we need a demand system that satisfies three properties. First, it must derive from individual consumer behavior, to enable us to measure welfare. Second, it must satisfy the typical shape properties of a valid demand system so we can compute optimal category prices. Third, it must be parsimonious so we can capture product differentiation without exceeding the limits of our store data in terms of degrees of freedom. We begin by specifying an economic model of individual choice behavior in a supermarket product category. We then derive the expected aggregate demand facing the category manager. Using the derived demand, we then model the category manager's pricing problem.

2.1. Utility and demand

We use the mixed logit consumer specification (McFadden and Train, 2000), which adds normally-distributed random coefficients to a standard logit choice model. The

3 By revealed preference, we know that profits must be at least as high under price discrimination since the non-discriminatory prices are in the feasible set.

4 Most of the theoretical 3rd-degree price discrimination literature looks at pricing for a single product across multiple market segments. However, retailers typically offer a product line in each market segment. Our analysis implies a layer of 2nd-degree price discrimination due to the selection of the product line (Cohen, 2000). We hold the product line decisions fixed when setting store-level prices so that the same goods are available in all stores.

random coefficients generate correlations in utilities for the various alternatives and, thus, relax the restrictive substitution patterns generated by the Independence of Irrelevant Alternatives (IIA) property of the logit model. A more general correlated additive error, such as the probit model, would avoid the IIA property directly and disentangle heterogeneity from simple non-IIA behavior at the consumer level. We use the mixed logit mainly due to its relative computational ease. However, McFadden and Train (2000) show that the mixture of normals with the type I error may be sufficiently flexible to approximate a broad set of parametric indirect utility functions, including the probit.

The discrete choice approach is parsimonious. Consumer preferences are projected onto a set of exogenous product attributes, which greatly reduces the dimension of the estimation problem. For industries with a large number of differentiated alternatives, correlations in product valuations may be characterized by heterogeneous tastes for the attributes. In many product markets, such as automobiles (Berry et al., 1995; Petrin, 2001; Sudhir, 2001a), researchers have readily-available measurable attributes to capture the underlying market segments in the category. In packaged goods product markets such as one would find in supermarkets, the attributes are largely intangible in nature and are reflected in the brand-specific intercepts (or intrinsic brand preferences) included in the model. Hence, we model the joint distribution of these brand preferences explicitly. We use a parsimonious factor-analytic approach for the covariance matrix of brand preferences, as is typically used with individual data (Elrod, 1988; Chintagunta, 1994; Elrod and Keane, 1995) as well as with aggregate data (Chintagunta et al., 2002).[5]

Formally, we assume that on a given shopping trip in week $t (t = 1, \dots, T)$ in store $s (s = 1, \dots, S)$, M_{ts} consumers each select one of J brands in the category or opt for the no-purchase alternative. The population of consumers in each store, s, is characterized by the vector of demographic variables, $\mathbf{D}_s$. Each brand j has attributes: $(\mathbf{x}_{jts}, \xi_{jts})$. The vector $\mathbf{x}$ includes dummy variables for promotion incidence and package size (e.g., ounces). The vector ξ encompasses the effects of unobserved (to the econometrician) weekly in-store product attributes, such as advertising, shelf-space and coupon availability that vary across store-weeks (Besanko et al., 1998; Nevo, 2001).[6] These unobserved factors generate deviations from the mean utility for a product across weeks and stores. In the estimation section below, we explain how these deviations from the mean might bias estimation. Finally, the variable p_{jts} denotes brand j's shelf-price in week t and store s.

5 Additional methods exist for identifying flexible substitution patterns using similar data with aggregate choices. For instance, Nevo (2001) samples consumer demographic profiles from the empirical joint distribution provided by the census at the MSA level. Berry et al. (1998) construct additional moments of the data-generating process by combining additional micro data with their aggregate data.

6 Since we estimate a full set of brand intercepts, we do not need to worry about unmeasured physical product attributes, as in Berry et al. (1995). We are nonetheless concerned with unobserved (to the econometrician) weekly in-store product-specific effects.

For a shopping trip during week t in store s, the conditional utility consumer h derives from purchasing product j is given by:

$$u_{hjts} = \alpha_{hjs} + \mathbf{x}_{jts}\beta_h + \theta_{hs}(Y_h - p_{jts}) + \xi_{jts} + \varepsilon_{hjts},$$
$$h = 1,\ldots,H, \quad j = 0,\ldots,J, \quad t = 1,\ldots,T, \quad s = 1,\ldots,S$$

where

$$\begin{bmatrix}\beta_h \\ \theta_{hs}\end{bmatrix} \sim N\left(\begin{bmatrix}\bar{\beta} \\ \bar{\theta} + D'_s\gamma\end{bmatrix}, \begin{bmatrix}\lambda_x & 0 \\ 0 & \lambda_p\end{bmatrix}\right),$$

and the vector $\boldsymbol{\alpha}_{hs} = (\alpha_{h1s},\ldots,\alpha_{hJs})'$ has the form

$$\boldsymbol{\alpha}_{hs} \sim N(\bar{\alpha} + D'_s\sigma, \Sigma).$$

The consumer-specific coefficient α_{hj} captures consumer h's intrinsic taste for brand j, β_h captures the taste for measured product attributes and θ_h is the marginal utility of income (valuation of expenditures outside the category). The coefficients $\bar{\beta}$ and $\bar{\theta}$ capture the mean tastes for observed attributes and mean marginal utility of income respectively across consumers in all stores. In the current context, income, Y_h, consists of the shopping budget for a trip during week t.[7] The parameters in the vector γ capture the deviations from the mean marginal utility of income across stores due to demographic differences. In practical terms, these deviations allow price elasticities to differ across stores. The parameters λ_x and λ_p capture the standard deviations of the β_h and θ_{hs} parameters across the population of consumers within a store. Similarly, the parameters $\bar{\alpha}_j$ capture the mean perceived intrinsic utility of brand j and the matrix Σ captures the covariance in these perceived values across consumers. The vector $\boldsymbol{\sigma}$ captures differences in mean perceived brand values across stores due to demographic differences. In practical terms, these deviations allow the size of the category (e.g., total inside share) to vary across stores.[8] The term ε_{hjts} is an i.i.d. draw from the type I extreme value distribution capturing a consumer's idiosyncratic utility for alternative j.

In theory, we could estimate all the elements of the $(J \times J)$ matrix Σ directly. In practice, as the number of brands grows, Σ becomes increasingly difficult to identify.

7 In the following analysis, we do not address formally how households allocate total income to their weekly shopping budgets.

8 With a long enough time series for each store, one could estimate a separate θ and α parameter for each store. This approach would require 166 parameters, which is infeasible with our 52-week sample. We focus on a 52-week window to avoid problems from potentially time-varying preferences. Our decomposition has additional potential benefits, such as the ability to forecast demand in a new store based on its characteristics.

Instead, we use the factor structure:

$$\Sigma = L\boldsymbol{\omega}\omega' L', \qquad \boldsymbol{\omega} \sim N(0, I).$$

One interpretation for this structure is that L is a $(J \times K)$ matrix of latent attributes for each of the J brands, and $\boldsymbol{\omega}$ is a $(K \times 1)$vector of consumer tastes for these attributes. The vector of mean brand perceptions, $\overline{\boldsymbol{\alpha}}$, and the matrix of latent attributes, L, consist of parameters to be estimated. In addition to its parsimony, this approach allows us to estimate standard errors for the latent attributes. In the current context, we assume $K = 2$. For identification purposes, we do the following (see Elrod, 1988):

1. The outside or "no purchase" option is located at the origin of the map (translational invariance).
2. One of the brands is located along the horizontal axis (rotational invariance).
3. We set the variances of $\boldsymbol{\omega}$ above to 1 in the estimation (scale invariance).

As noted above, the formulation allows for an outside good, "no purchase", the utility of which is given by

$$u_{ht0s} = \alpha_{h0s} + \theta_{hs} Y_h + \varepsilon_{ht0s}.$$

In the current context, this alternative represents the allocation of the shopping budget, Y_h, to other goods in the store outside the category. For practical reasons, this outside good is important for the retailer pricing exercise. In the absence of this alternative, the total category size would be invariant to the prices of all brands increasing or decreasing by the same amount. Hence, allowing for the outside good allows the category sales to be influenced by the prices of the inside goods. For identification purposes, α_{h0s} is normalized to zero.

As is now the convention in the literature, we simplify our notation by re-writing the consumer's indirect utility in terms of mean tastes and deviations from the mean:

$$u_{hjts} = \delta_{jts} + \mu_{hjts} + \varepsilon_{hjts}, \tag{1}$$

where $\delta_{jts} = \left(\overline{\alpha}_j + D'_s\sigma\right) + x_{jts}\overline{\beta} - \left(\overline{\theta} - D'_s\gamma\right)p_{jts} + \xi_{jts}$ is common across consumers and $\mu_{hjts} = v'_{hx}\mathbf{x}_{jts}\lambda_x - \boldsymbol{v}_{hp}p_{jts}\lambda_p + L\omega_h$ is consumer-specific, where $\boldsymbol{v} = \left(v_{hx}, v_{hp}, \omega_h\right)'$ is a vector of standard normal deviates.[9] An advantage of this mixture of the normally-distributed tastes with the extreme value disturbance is that we can integrate out the latter analytically. The unconditional probability q_{jts} that a consumer chooses a particular product j in week t and store s, after integrating

9 Since $\theta_h Y_h$ is common to all the alternatives, it is not identified in the demand system below. So we remove it from the equations.

across the joint distribution of $\boldsymbol{\nu}$, has the following form:

$$q_{jts} = \int \cdots \int_{-\infty}^{\infty} \frac{\exp(\delta_{jts} + \mu_{hjts})}{1 + \sum_{i=1}^{J} \exp(\delta_{its} + \mu_{hits})} \phi(\boldsymbol{\nu}) \partial \boldsymbol{\nu}, \tag{2}$$
$$h = 1, \ldots, H, \quad j = 0, \ldots, J, \quad t = 1, \ldots, T, \quad s = 1, \ldots, S$$

where $\phi(\cdot)$ is the pdf of a multivariate standard normal. From the store manager's perspective, (2) represents the share of consumers entering the store in week t that purchase a unit of product j. Thus, the manager's expected demand for product j in week t and store s is :

$$Q_{jts} = q_{jts} M_{ts}. \tag{3}$$

The random coefficients add flexibility to the empirical model by mitigating the effects of the IIA property, which could manifest itself into our empirical analysis in several ways. First, it can be shown that a homogeneous logit model generates aggregate cross-elasticities that are driven by market shares. For instance, products with similar market shares are predicted to be close substitutes. On the supply-side, the cross-elasticities also restrict the implied retailer margins. Multiproduct firms are restricted to set a uniform margin for each of the products in their line (Besanko et al., 2001a). In our category pricing simulations, retail mark-ups over wholesale prices would be equal across all the products. Adding random coefficients relaxes these restrictions. Note that the inclusion of random brand intercepts also allows for flexible substitution patterns between purchase and no-purchase decisions.

Since we do not observe competitors' prices in our data, we cannot model competition explicitly. Instead, we assume that any local market power is captured, on average, by proximity to competitors and include variables reflecting these distances in the vector $\mathbf{D}_s$ described previously. To the best of our knowledge, no study has modeled store competition explicitly in determining aggregate demand for a retailer. In modeling demand for yogurt, Berto Villas-Boas (2001) treats the same brands in different stores as substitutes. However, she finds extremely small cross-price elasticities across stores, suggesting that store competition has little impact within the yogurt category. Most applications of store-level data treat retailers as local monopolists (e.g., Slade, 1995; Besanko et al., 1998). Slade justifies her assumption on phone interviews with store managers in a given market who claim that consumers do not shop across stores on a product-by-product basis. We also conducted telephone interviews with Chicago area store managers and our findings were consistent with this claim. Stores do condition on their competitors' actions in a limited way by collecting a weekly sample of half a dozen SKUs from the local competitors' entire store offerings. However, this behavior seems more consistent with competition on overall offerings rather than on a category-by-category basis.[10] For instance, Chevalier et al. (2000) find that prices for items exhibiting holiday or seasonal demand peaks tend to be priced in a manner consistent with loss-leader competitive pricing (Lal and Matutes, 1994). Therefore, we try to limit our focus to product categories that are less likely to drive overall store traffic.[11] Furthermore, Dreze et al. (1994b) run a pricing experiment using a random subset of the same DFF stores. After increasing prices by 9% across 26 categories (33% of total store volume) for 16 weeks, they find no significant effect on store traffic (the weekly customer count) and no significant effect on dollar sales in the non-test categories. Given these findings, we do not expect price changes in any single category to generate substantial losses or gains in total store traffic. We discuss this assumption further in the data section below.

2.2. *Measuring consumer welfare*

One of our main objectives in using a structural demand system is the ability to measure the change in consumer welfare associated with altering a chain's pricing policy. An attractive feature of the discrete choice model is the ability to compute consumer welfare explicitly. A popular measure for welfare in such contexts is the Hicksian, or compensating, variation, which captures the dollar amount by which consumers would need to be compensated to maintain the same level of utility after the change in pricing policy (e.g., Trajtenberg, 1989; Nevo, 2000; Petrin, 2001). In a store s, we denote an individual h's utility net of the extreme value taste shock as V_{hs} (expected utility) and their marginal utility of income as θ_{hs}. Suppose a zone-pricing policy is introduced that changes consumer valuations for each alternative from V_{hs}^{chain} to V_{hs}^{zone}. As derived in Small and Rosen (1981), assuming individual marginal value of income is held constant, individual h's associated change in welfare can be computed as:

$$CV_{hs} = \frac{\log\left(\sum_{j=0}^{J} \exp\left(V_{jhs}^{zone}\right)\right) - \log\left(\sum_{j=0}^{J} \exp\left(V_{jhs}^{chain}\right)\right)}{\theta_{hs}}. \tag{4}$$

The numerator of (4) captures the expected change in utility, the difference between the expected maximized utility under the two pricing policies. Dividing through by the marginal utility of income, θ_{hs}, makes this change money-metric. The value CV_{hs} thus measures the dollar amount by which consumer h must be compensated to be as well off under zone pricing as under chain-wide pricing. Integrating across the

10 Stores collect a "full book" of about 600 to 1000 prices from local competitors annually. This practice is not likely to generate inter-store competition at the category level and weekly frequency we consider in our analysis.

11 We provide empirical support below, when we measure the impact of zone pricing.

distribution of consumer preferences in the store, we can compute the expected aggregate change in consumer welfare:

$$\Delta W_s = M_s \int_{-\infty}^{\infty} \cdots \int CV_{hs}\phi(v)\partial v. \tag{5}$$

The value ΔW_s captures the total compensation required for all consumers in a store-week. In the next section, we model the supermarket category manager's pricing decision. Using variable profits as the measure of the manager's valuation, we are able to compare the dollar value of gains of various pricing policies both to the supermarket and to consumers. Similarly, we can compute the change in customer value in going from zone to store level pricing.

In computing consumer welfare, we find yet another role for controlling for heterogeneity in tastes across consumers within a store. In the absence of such heterogeneity (e.g., consumer heterogeneity is captured entirely through the extreme value error term), generating welfare gains due to a change in pricing policy requires category expansion. To see this point, suppose moving from chain to zone pricing increases welfare:

$$\begin{gathered}
\log\left(\sum_{j=0}^{J} \exp\left(V_{js}^{zone}\right)\right) > \log\left(\sum_{j=0}^{J} \exp\left(V_{js}^{chain}\right)\right) \\
\sum_{j=0}^{J} \exp\left(V_{js}^{zone}\right) > \sum_{j=0}^{J} \exp\left(V_{js}^{chain}\right) \\
\sum_{j=0}^{J} \exp\left(V_{js}^{zone}\right) + 1 > \sum_{j=0}^{J} \exp\left(V_{js}^{chain}\right) + 1 \\
\frac{1}{\sum_{j=0}^{J} \exp\left(V_{js}^{zone}\right) + 1} < \frac{1}{\sum_{j=0}^{J} \exp\left(V_{js}^{chain}\right) + 1}. \\
q_0^{zone} < q_0^{chain}
\end{gathered}$$

Thus, welfare gains necessitate category expansion, or decreases in the share of the outside good, for the case of the homogeneous logit. In this particular instance, it is easy to show category expansion is also a sufficient condition. However, once we add the random coefficients, one can generate numerical examples in which welfare may rise irrespective of category expansion (see Chiang and Spatt, 1982, for a similar outcome). This result has intuitive appeal for frequently-purchased goods, as one finds in supermarkets, since consumer gains do not strictly require the conversion of non-purchasers. The aggregate data limitations do not permit us to place much structure on the nature of non-purchasers. Thus, non-purchase is an average across consumers that may indeed prefer expenditure in other categories as well as consumers that may occasionally hold inventories of the inside goods, temporarily lowering their willingness-to-pay. Category expansion could be a misleading necessary condition for the latter type of consumer.

2.3. *Category pricing*

We now describe our model of retail behavior. Our data comprise stores from a single retail chain in a large metropolitan area. We assume that each week the retailer jointly sets the profit-maximizing prices, p_j of each of the J brands in a category, given wholesale prices for each of the J brands (see, for example, Sudhir, 2000b; Kadiyali et al., 2000). Based on our phone conversations with local chain managers, we believe that weekly price decisions are made by category managers rather than by a store-wide manager.[12] In contrast, most promotional decisions (newspaper features, in-aisle displays etc.) are determined at the chain-level. While promotions are funded almost entirely by manufacturers, the timing and format are determined by the retailer. Typically, the promotional calendar is determined in advance so that category pricing decisions are made conditional on the promotion. Therefore, we treat the promotion level as exogenous to category pricing.[13]

We also assume that the retailer's variable costs consist solely of wholesale prices, w_j. We treat all store and/or category-related overhead as fixed costs, F_t. Thus, in week t, a retailer[14] solves the following optimization problem:

$$\max_{\{p_j\}_{j=1}^{J}} \Pi = \sum_{j=1}^{J} (p_{jt} - w_{jt})Q_{jt} - F_t.$$

Using (3) above, the first-order condition for a typical brand i is:

$$\sum_{j=1}^{J} (p_{jt} - w_{jt}) \frac{\partial Q_{jt}}{\partial p_{it}} + Q_{it} = 0.$$

12 The retailer's category definition may differ from that of academic research. The latter typically uses the Nielsen and IRI definitions, the two traditional suppliers of comparable scanner databases.

13 Our phone conversations revealed that category managers may in fact request additional promotional funds if they feel the performance of the category or a specific brand therein is sluggish. However, the incidence of such endogenous (to the category manager) promotions are quite unusual.

14 To simplify notation, we drop the subscript for the retailer and the category. In our empirical work, a retailer may be the store-manager, the zone manager or the chain manager depending on the context.

We re-write the system of first-order conditions for brands $1, \ldots, J$ in matrix form as:

$$\Omega(\mathbf{p} - \mathbf{w}) + \mathbf{Q} = 0, \tag{6}$$

where

$$p - \mathbf{w} \equiv \begin{bmatrix} p_{1t} - w_{1t} \\ \vdots \\ p_{Jt} - w_{Jt} \end{bmatrix}_{J \times 1}$$

$$\Omega_{jk} = \begin{cases} \dfrac{\partial Q_{jt}}{\partial p_{kt}}, j = k \\ \dfrac{\partial Q_{jt}}{\partial p_{kt}}, j \neq k \end{cases}_{J \times J}$$

$$\mathbf{Q} = \begin{bmatrix} Q_{1t} \\ \vdots \\ Q_{Jt} \end{bmatrix}_{J \times 1}.$$

This represents a system of J equations, one for each brand. The optimal set of prices for the retailer are determined by solving:

$$\mathbf{p} = \mathbf{w} - \Omega^{-1}\mathbf{q}, \tag{7}$$

where $\Omega^{-1}\mathbf{q}$ is the retail mark-up. By checking the second order sufficient conditions, one can verify that the solution to (7) represents an optimum for the retailer. The actual value of Q_{jt} depends on the level of aggregation considered. For the store-level problem, this will take the form (3), where M_t is weekly store traffic. However, for a zone pricing problem, Q_{jt} would be obtained by integrating across the store-level demand for each of the stores in the zone in week t. Similarly, chain-level pricing would involve integrating across all the store-level demand curves.

As we mentioned in the demand section, the assumption of homogeneous tastes leads to restrictive pricing behavior by retailers. When consumer tastes for attributes are homogeneous, the optimal retail prices satisfying (6) become:

$$p_{jt} = w_{jt} + \frac{1}{\theta q_{0t}},$$

where $q_{0t} = 1 - \sum_{j=1}^{J} q_{jt}$ is the share of the no-purchase alternative. Therefore, in the absence of consumer heterogeneity, a retailer's mark-ups over the wholesale prices are the same for all the products carried. This property property of the homogeneous logit model is not consistent with our data in which margins vary across alternatives.

3. Estimation

We now outline the estimation procedure for the aggregate mixed logit model described above. Since one of our objectives in this analysis is the determination of the level of aggregation at which stores determine prices, we estimate demand alone and do not consider the supply-side in the estimation of demand parameters. This ensures that our demand-side estimates are not subject to specification error from assuming a specific category management model by retailers. Our estimation methodology is quite similar to that used by Berry et al. (1995). Therefore, we refer the more interested reader to Berry et al. (1995) for a more technical description and to Nevo (2001) for a more practical discussion of the implementation of the methodology.

A primary concern in empirical papers using similar discrete choice models is the potential for estimation bias due to correlation between prices and the unobserved product attribute, ξ. Using weekly store-level data, the issue lies in unmeasured store-specific covariates that influence demand and also shift prices at a weekly frequency. Even after including a full set of alternative-specific intercepts, several papers have documented evidence of an estimation bias in models that do not control for this problem using weekly supermarket data (Besanko et al., 1998; Chintagunta, 2002; Villas-Boas and Winer, 1999; Villas-Boas and Zhao, 2001). For instance, we do not observe shelf-space; however, increasing shelf-space allocation typically incurs costs that raise prices, such as allocation fees and opportunity costs. At the same time, it is well known that shelf-space influences consumer brand choices (Dreze et al., 1994a). While characterizing the precise nature of such measurement error in our data is beyond the scope of the paper, we use standard instrumental variable techniques to avoid estimation biases.

We begin by partitioning the observed marketing variables for each brand j in store s and week t as $X_{jts} = [x_{jts}, p_{jts}]$, where by assumption $E(x_{jts}\xi_{jts} \mid x_{jts}) = 0$ and $E(p_{jts}\xi_{jts} \mid p_{jts}) \neq 0$. Following Berry (1994), we invert (2) to recover the mean utilities $\delta_{jts}(\Theta)$as functions of a vector of model parameters, Θ, and set up the estimation procedure in terms of δ_{jts}. Recall from equation (1), δ_{jts} captures the mean utility of brand j in store s during week t, where this mean utility incorporates the mean tastes for observed attributes, x_{jts}, as well as the unobserved attributes, ξ_{jts}. Thus, we compute values of δ_{jts} that fit the predicted shares in (2) exactly to the observed market shares. Note that the treatment of heterogeneity in the model introduces a multidimensional integral in the share equation (2). We evaluate these integrals using Monte Carlo simulation. Also, since the inverse of (2) does not have a simple analytical form, we use the contraction-mapping of Berry et al. (1995). The advantage of using δ_{jts} for estimation is that the prediction error,

$\delta_{jts} - \overline{\alpha_j} - X_{jts}\overline{\beta} + \overline{\theta}_s p_{jts}$, is simply the unobserved product characteristic, ξ_{jts}. Recall that ξ_{jts} enters the mean utility linearly, $\delta_{jts}(\Theta) = (\overline{\alpha}_j + D_s'\delta) + \mathbf{x}_{jts}\overline{\beta} - (\overline{\theta} - D'_s\gamma)p_{jts} + \xi_{jts}$, which facilitates instrumentation.

Defining $\xi_{ts} = (\xi_{1ts}, \ldots, \xi_{Jts})$, our formal data-generating process is described by the conditional mean-independence assumptions, $E(\xi_{ts} \otimes \mathbf{Z}_{ts} | \mathbf{Z}_{ts}) = 0$, where $\mathbf{Z}_{ts}$ is a vector of exogenous instruments. Note that we impose fewer covariance restrictions than previous papers in this area by assuming $E(\xi_{ts}\xi'_{ts}|\mathbf{Z}_{ts}) = \Omega$ a finite $(J \times J)$matrix and $E(\xi_{ts}\xi'_{\tau s}) = 0,\ t \neq \tau$. By modeling the contemporaneous covariances of the unobserved attributes, we rely on the large number of weeks and stores for identification, rather than on the size of the cross-section of products themselves. We fit these moments to our data using the generalized method of moments (GMM) procedure of Hansen (1982). As explained above, the computation of the moment conditions requires simulating the share equations. McFadden (1989) and Pakes and Pollard (1989) both show that this method of simulated moments (MSM) produces consistent estimates. However, the efficiency of these estimates is reduced due to simulation error. Only with sufficiently many simulation draws can one reach asymptotic efficiency with MSM. We use 30 draws and assume this number is sufficient to eliminate any noticeable simulation noise. Alternatively, one could implement variance-reducing simulation methods as in Berry et al. (1995).

4. Data

We use data from DFF, which is the second largest supermarket chain in the Chicago metropolitan area. DFF operates close to 100 stores in the Chicago area. Our data consist of weekly sales, prices, promotions, and profit margins at the individual UPC-level for 83 of these stores during the 52-week period of 1992. In the current analysis, we look at the liquid laundry detergent and refrigerated orange juice categories. We define products as brand-size combinations (e.g., a 64 oz Tropicana Premium vs. a 96 oz). The promotion variable is an indicator for whether the given product had an in-aisle display or newspaper feature that week. As mentioned previously, the promotion decision is assumed to be exogenous to category management within a store or zone. Note that in some instances we aggregated UPCs with the same brand-size combination (e.g., 64 oz Tropicana with and without pulp). To avoid aggregation biases, we only aggregated such products when their correlation in prices was greater than 0.9. We present the descriptive statistics for the products included in the analysis for each of the respective categories in Table 1. These data consist of means across store-weeks. The column labeled "unit share" corresponds to the conditional shares, or share of unit sales, for each brand. The actual market shares used for estimation are computed as total brand sales divided by the total weekly store traffic. Effectively, our model implies that each time consumers visit the store, they either purchase a unit of one of the alternatives within the category of interest or they elect not to purchase.

Table 1. Descriptive statistics.

Category	Brand	Size	Share	Price	Cost	Prom
Refrigerated orange juice	Minute Maid	64 oz	21.2%	2.24	1.69	0.37
	Minute Maid	96 oz	4.1%	4.09	1.99	0.31
	Dominicks	64 oz	24.4%	1.65	1.15	0.39
	Tropicana Premium	64 oz	20.5%	2.66	1.88	0.39
	Tropicana SB	64 oz	17.7%	2.37	1.62	0.29
	Tropicana Premium	96 oz	7.7%	4.51	2.18	0.14
	Florida	64 oz	4.3%	2.18	1.90	0.22
Laundry detergent	Surf	64 oz	6.2%	4.09	3.01	0.27
	Wisk	128 oz	7.0%	8.10	3.31	0.13
	Wisk	64 oz	14.1%	4.14	3.53	0.17
	All	64 oz	12.8%	3.11	2.41	0.24
	All	128 oz	10.9%	5.72	2.19	0.13
	Cheer	64 oz	6.3%	4.20	3.62	0.16
	Cheer	128 oz	5.3%	8.20	3.42	0.25
	Tide	128 oz	18.9%	8.30	3.52	0.41
	Tide	64 oz	18.4%	4.38	3.79	0.25

The computation of the market shares assumes single-unit purchases at the individual consumer trip level. In some product categories, this assumption may be inappropriate. Consumers may purchase varying quanitites of a single brand or assortments consisting of multiple brands and varying quantities of each. A small marketing literature has explored the benefits of models that account for consumers' quantity decisions for a single brand (Chiang, 1991; Chintagunta, 1993) and for consumers' assortment decisions within a category (Dubé, 2001; Kim et al., 2001). Assuming single-unit purchases in such categories could understate the own and cross-price elasticities of demand. In the context of pricing, the model would tend to overstate the extent to which prices could be raised above costs profitably. For the categories used in the current analysis, we believe the single-unit purchase assumption is not overly restrictive. We verify this assumption using a household scanner panel database for 2108 households in Denver between January 1993 and May 1995. For laundry detergent, fewer than 2% of the trips during which an item was purchased involve the purchase of multiple product alternatives and fewer than 10% involve the purchase of more than a single unit of any alternative. For refrigerated orange juice, fewer than 1% of the purchase trips in the category involve the purchase of multiple product alternatives and fewer than 10% involve the purchase of more than a single unit of any of the alternatives.

We supplement our store data with an extensive set of descriptive variables, from Spectra (see Hoch et al., 1995), to characterize the underlying consumer base and local competition associated with each store. ZIP code level demographic data was obtained from the 1990 census. To capture heterogeneity across stores we include the following demographic variables: INCOME (log median income), AGE60 (% of population over age 60), ETHNIC (% of population that are Black or Hispanic) and

Table 2. Demographic and competitive variables.

Variable	Mean	Std Dev.	Minimum	Maximum
INCOME (log)	10.6	0.28	9.9	11.2
AGE60	17%	6%	6%	31%
ETHNIC	15%	19%	2%	99%
HVAL	147.3	46.2	64.3	267.4
SHOPINDX	74%	24%	0%	99%
JEWELDIST	1.29 (mi)	0.86	0.06	3.96
EDLPDIST	5.03 (mi)	3.48	0.13	17.85

HVAL (mean household value). We also include SHOPINDX (ability to shop–% of population with car and single family house) to capture the relative ability of local consumers to travel. The two competitive variables used in the study are distance from the nearest Jewel (the largest supermarket in the area), JEWEL, and minimum of the distance from the nearest Cubfoods and Omni, EDLP, (the two main EDLP operations). Our preliminary work also included variables on competitor volume but these had limited explanatory power and were dropped in the subsequent stages of analysis.

Summary statistics for the demographic and competitive variables are provided in Table 2. We find considerable variation in the demographic and competitive characteristics across stores. For example, DFF stores cater to market areas with Black and Hispanic representation ranging from 2 to 99% of the population. In terms of consumer wealth, income levels range from $19,000 to over $75,000 (note we report INCOME in logs). Similarly, average house values range from $64,000 to over $267,000. In terms of competition, some stores are located right next to rival supermarkets. Others locate over 4 miles from the nearest Jewel and 18 miles from the nearest EDLP store. As we discuss below, we expect these differences to generate noticeable differences in the levels and price-sensitivities of demand across stores. Since the estimation of store-specific parameters would generate an unmanageable number of parameters for the given data sets, we have some assurance that the demographics do a reasonable job explaining store-specific differences.

4.1. Zones

Our data set contains an index that groups stores into 16 pricing zones. In Figure 1, we plot the stores on the Chicago area map, labeling each according to its zone affiliation from 1 to 16. In practice, the chain does not always appear to respect the specified zones in its weekly pricing decision. Looking across brands, we observe that many items appear to use a coarser zone definition. For instance, prices of small share items often have a uniform price across stores. Similarly, in some categories, prices may only reflect three or four zones, rather than the full 16 that we observe later in the data. Other studies that have used this data (for example Hoch et al.,

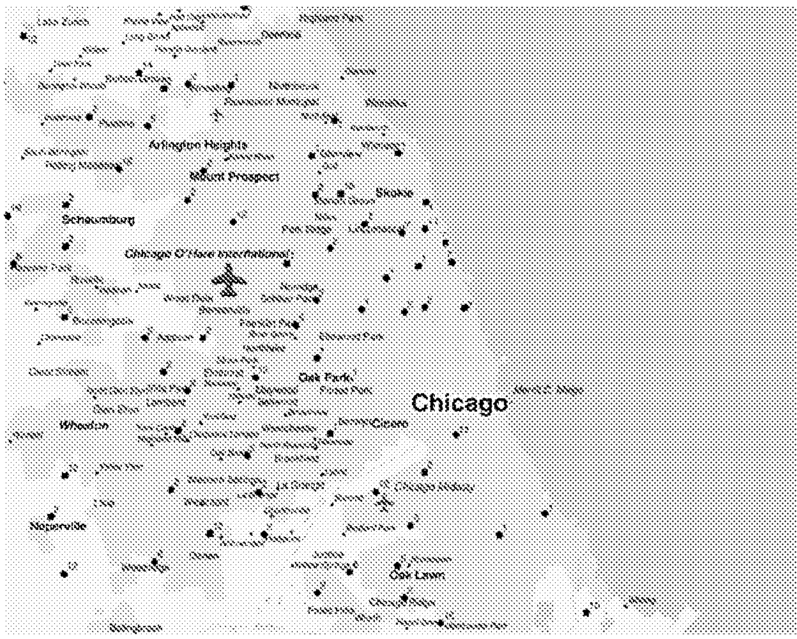

Figure 1. 16-zone configuration used by DFF.

1995) also suggest that the actual number of zones might be fewer than those provided in the data. We investigate this issue by looking at the prices of several brands across randomly selected weeks. Consistent with Hoch et al., we find only three levels of prices in the early years of the total available data (e.g., 1989–1990). However, the number of zones increases over time. By the time of our sample, 1992, we begin to observe substantially more prices for many large-share products in any given week. For some of the large share brands, we often observe between nine and sixteen different shelf prices across stores within a given week.[15]

To illustrate the degree of price-dispersion across stores in 1992, Figure 2 presents the distribution of prices for 128 oz Tide laundry detergent for a given week. Looking across stores, we see prices ranging from as low as \$3.80 to as high as \$4.90. However, we note that most of this price dispersion is driven by the reported zone structure. In the refrigerated orange juice category, the price of Minute Maid (64 oz) is 38% higher in the highest price zone, compared to the stores that fall in the lowest price zone. Similarly, price dispersion is much lower for stores within a zone than across zones. For liquid laundry detergent, the average range in prices for a given

15 Singh (2002) discusses two possibile explanations for the increase in zones over time: either the chain is facing a changing competitve environment (entry of mass-merchandisers and club stores) or Dominicks management is varying its pricing policy.

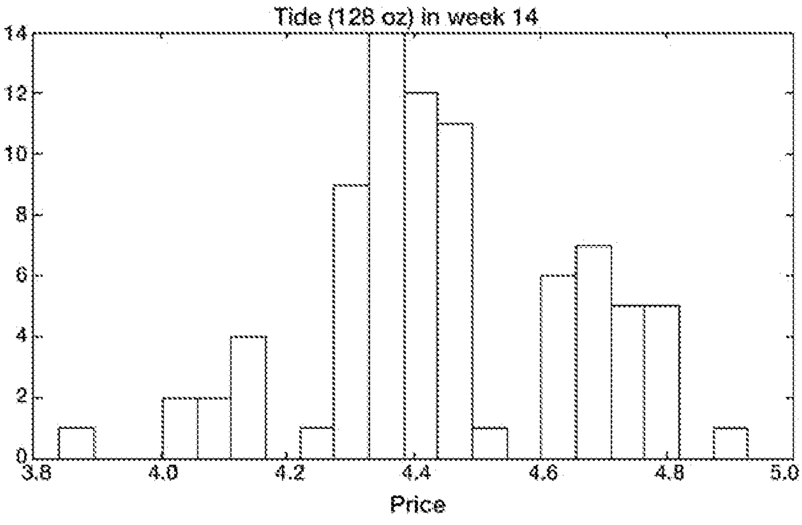

Figure 2. Distribution of shelf prices across stores for 128 oz Tide in a given week.

brand across stores in a given week is 81 cents. By contrast, the average range within a Dominicks-classified zone is about 16 cents. Given that the mean price in the category is \$5.58, the average weekly price range across stores overall is 14% of the mean price, versus only 3% within a zone. Thus, we are reasonably confident that the reported zone configuration reflects the pricing policies of Dominicks in 1992, the year used in our analysis.

Now that we understand the zone-pricing, we need to determine whether the zones themselves are motivated by price discrimination or alternative sources such as cost differences, consumer search behavior or competition. An important feature of our data is the ability to disentangle price discrimination and differences in cost. In general, wholesale prices are virtually identical across stores in a given week. On average, the standard deviation of wholesale prices across stores in a given week is 0.008. Unlike previous work (e.g., Shepard, 1991) we can easily rule out explanations for price variation based on wholesale costs.[16]

Having ruled out costs, we now classify our store characteristics to help us distinguish amongst the remaining explanations. We classify the store variables into three categories: willingness-to-pay, consumer search and competition. The demographic variables are used to capture consumer heterogeneity in tastes across stores. The extent to which these factors influence demand at a store is attributed to price discrimination. The Spectra measure SHOPINDX measures the ability of a

16 One explanation that we have not ruled out is differences in the opportunity cost of shelf space in areas with expensive real estate. In areas with high property values and/or high property taxes, stores may have more rigid capacity constraints that could affect pricing.

Table 3. Demographics and competitive variables by zone.

		Demographics						
Zone	Price ($)	INCOME (log $)	AGE60 (%)	ETHNIC (%)	SHOPINDX (%)	HVAL ($'000)	JEWEL (miles)	EDLP (miles)
1	3.48	10.56	0.20	0.26	0.62	166.52	1.29	7.24
2	3.27	10.66	0.17	0.11	0.79	147.37	1.19	4.20
3	3.38	10.62	0.18	0.03	0.89	143.83	1.04	2.10
4	3.28	10.80	0.21	0.05	0.80	160.00	2.47	1.63
5	3.26	10.59	0.22	0.07	0.83	111.59	1.76	2.80
6	3.20	10.71	0.14	0.08	0.87	135.32	1.26	1.14
7	3.48	10.43	0.21	0.23	0.42	190.35	1.35	8.51
8	3.27	10.56	0.13	0.15	0.88	121.19	0.34	11.65
10	3.46	10.57	0.26	0.15	0.75	116.41	1.83	7.28
11	3.38	10.10	0.15	0.46	0.25	97.37	1.58	9.53
12	3.26	10.74	0.16	0.14	0.84	152.75	0.93	4.89
13	3.26	10.72	0.09	0.11	0.94	151.07	3.91	6.68
14	3.26	10.88	0.09	0.07	0.76	179.07	1.43	3.10
15	3.25	10.57	0.14	0.21	0.88	100.39	1.80	4.41
16	3.19	10.78	0.06	0.08	0.81	139.06	1.72	2.60
Mean	3.31	10.62	0.16	0.15	0.76	140.82	1.59	5.18

consumer to shop and, thus, to search for the "best" price in the local market. Finally, proximity to competitors provide a crude measure of the level of local competition. The extent to which each of these variables shifts demand, and thus prices, within a zone will help us determine whether or not Dominicks is currently price discriminating. In Table 3 we report the average laundry detergent prices and store characteristics by zone. These data demonstrate the ability of our store characteristics to explain some of the observed price differences across zones. For instance, the highest prices are in zone 7, where we observe the highest household values as well as fairly high search costs. In contrast, zone 16 has the lowest prices and exhibits fairly high incomes with very few elderly or ethnic households and a higher than average shopping ability. Note also that zone 6 has the closest EDLP store and, at the same time, has very low prices. This table simply allows us to conclude that each of the store characteristics may be partially responsible for price dispersion. Ultimately, we will need to look at the marginal effect of these variables on demand to assess which explanation is most relevant.

4.2. Instruments

An important component of our analysis is the ability to instrument prices and control for potentially unmeasured product characteristics. In selecting a good instrument, we first look at the nature of price variation in our data. We find that the median R^2 across products for a regression of price on week dummy variables is roughly 72%, in the laundry detergent category, and 77% in refrigerated orange juice. Similarly, the median R^2 across products for a regression of price on store dummy variables explains roughly 11% of the price variation, in laundry detergent, and 6% in orange juice.

An unusual feature of our data is the availability of wholesale prices. These wholesale prices are specific to each brand in the category. We treat these data as exogenous to store-level demand in this chain.[17] We also show below that these wholesale brand prices correlate well with retail brand prices. As they are exogenous to store-level demand and hence unobserved product characteristics in the store and as they are correlated with retail prices, we use wholesale prices as instruments for in-store retail prices. In other contexts, marketers have modeled the vertical channel to capture the strategic interaction between retail and wholesale prices using a logit demand model (e.g., Besanko et al., 1998; Sudhir, 2001b; Villas-Boas and Zhao, 2001). Typically, these studies do not have access to retail margins and, thus, use the channel structure to help identify a time-varying wholesale price (see, for example, Berto Villas-Boas, 2001). Kadiyali et al. (2000) use information on retail margins. However, their objective was to identify the nature of interactions between manufacturers and a single retailer.

Although not reported, a pooled regression of shelf prices on wholesale prices alone gives an R^2 of 0.71 (refrigerated juice) and 0.74 (laundry detergent). Running the regression by SKU, we find that the wholesale prices do a much better job of explaining the larger-share brands of detergent (R^2 of about 0.6 on average) than the smaller brands (R^2 of about 0.1 on average). For refrigerated juice, the wholesale prices alone explain about 25% of the variation, on average, regardless of share. Introducing the additional exogenous covariates used in the instrument matrix for the GMM procedure tends to explain an additional 10% of the variation in the prices. Given the reasonably strong explanatory power of these instruments, we are able to identify the structural demand parameters without using supply-side moments. In previous research for which good instruments have not been available, researchers have relied on supply-side moment restrictions to help identify the price parameter. To illustrate the ability of our wholesale price data to explain price variation, we plot the time-series for prices and wholesale prices of 128 oz Tide Laundry Detergent for one of the stores in Figure 3. The figure shows that retail price movements generally coincide with changes in wholesale prices. Hence, these instruments capture the most important aspect of product-specific price variation which is across weeks. However, it does appear that retail passthrough rates exceed 100% in certain weeks. In the following section, we briefly discuss the implications of not instrumenting for prices, which generates noticeable biases in both the model parameters and in the predicted zone prices.

17 Note that DFF only accounts for 25% of Chicago areas supermarket sales. Thus, it is unlikely that a marginal change in price in one of DFF's 16 zones would shift aggregate Chicago demand sufficiently to impact the market wholesale price.

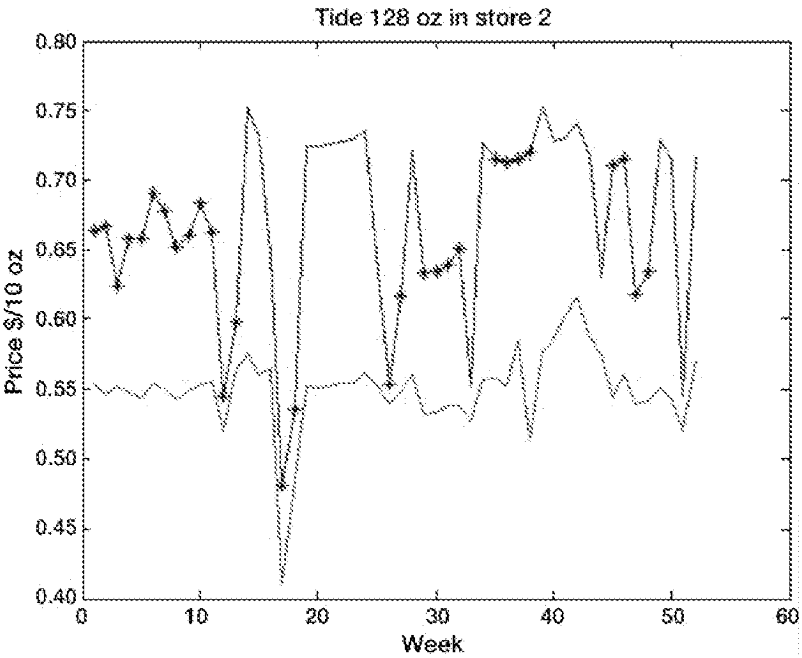

Figure 3. Shelf prices and wholesale prices (* indicates promotion week).

5. Results

We now present the estimated demand parameters for each of the two categories. These parameters are reported in Tables 4 and 5 for laundry detergents and orange juice respectively. Rather than reporting the implied brand correlation matrices, Σ, we plot the brand "locations" using the latent factors as a coordinate system. These maps are in Figures 4 and 5 for the two categories. We also report the elasticities of store characteristics on category size (probability of purchase) in Table 6.

Our first task is to motivate the endogeneity problem. In each category, we estimate a homogeneous logit demand system using OLS and compare the estimated mean price response to that of a system estimated using two stage least squares (see Besanko et al., 1998 for technical details). In refrigerated orange juice, the mean price response is -9.95 using OLS versus -14.65 using instrumental variables, a difference of 47%. In laundry detergent, the mean price response is -10.87 using OLS versus -12.13 using instrumental variables, a difference of 12%. A Hausman test of the hypothesis that the mean price response is statistically equivalent under OLS and two-stage least squares yields a statistic of 289.14 in orange juice and 59.07 in laundry detergent. In both cases, we easily reject the null hypothesis, finding the

Table 4. Demand for laundry detergent.

	Parameters	S.E.
Surf	– 6.273	2.654
Wisk	– 7.216	2.786
All	– 5.316	2.554
Cheer	– 3.305	2.550
Tide	– 1.866	2.541
price	– 7.843	4.082
s.d. price	0.348	0.393
promo	0.749	0.068
price* promo	– 1.106	0.112
64-oz	0.742	0.008
holiday	0.227	0.026
income	0.412	0.262
age60	– 0.907	0.420
ethnic	0.064	0.207
shopindx	0.436	0.255
hval	– 0.009	0.001
Jewel	– 0.002	0.029
EDLP	– 0.010	0.008
price*income	– 0.713	0.421
price*age60	3.296	0.672
price*ethnic	– 0.641	0.324
price*shopindx	– 0.273	0.410
price*hval	0.020	0.002
price*Jewel	0.050	0.046
price*EDLP	0.053	0.013

instrumental variables estimate to be significantly larger in magnitude than OLS. This result implies that OLS yields less elastic estimates of demand than two-stage least squares, as we hypothesized above. To illustrate the impact of this difference in the orange juice category, we find the average own-price elasticity of demand to be – 3.55 with OLS versus – 5.22 with two-stage least squares. Moreover, the implied optimal category manager's mark-up would be 83 cents under OLS versus only 57 cents under two-stage least squares. These result suggest that ignoring the potential endogeneity of unmeasured brand characteristics could lead a manager to overestimate consumer willingness-to-pay and, ultimately, to set prices too high.

We now examine the results from the heterogeneous logit demand systems presented in the previous sections. Before describing the empirical results, we first explain what the parameters in Tables 4 and 5 mean. We use Table 4 (laundry detergent) as an illustration. The first five parameters correspond to the mean intrinsic preferences, $\overline{\alpha_j}$, for the J brands included in the category. This is followed by the mean price effect, $\overline{\theta}$, and the standard deviation of the distribution of price sensitivities across consumers, λ_p. This is followed by the effects of promotion and

Table 5. Demand for refrigerated orange juice.

	Parameters	S.E.
Minute Maid	10.858	1.823
Dominicks	9.802	1.924
Tropicana Premium	11.953	1.817
Tropicana SB	11.008	1.814
Florida	8.724	1.969
price	− 50.389	4.443
s.d. price	0.033	0.874
promo	1.632	0.079
price* promo	− 3.694	0.211
96-oz size	− 0.485	0.011
holiday	0.198	0.031
income	− 0.921	0.186
age60	− 1.933	0.284
ethnic	0.375	0.147
shopindx	− 0.062	0.170
hval	− 0.015	0.001
Jewel	0.012	0.020
EDLP	− 0.029	0.006
price*income	2.950	0.459
price*age60	8.096	0.693
price*ethnic	− 0.436	0.346
price*shopindx	− 0.004	0.419
price*hval	0.049	0.002
price*Jewel	− 0.062	0.049
price* EDLP	0.120	0.013

Table 6. Marginal effects of store characteristics on category size (purchase probability).

Category	INCOME	AGE60	ETHNIC	SHOPINDX	HVAL	JEWEL	EDLP
Laundry detergent	− 0.107	0.126	− 0.018	0.160	0.301	0.024	0.067
Refrigerated OJ	0.751	0.109	0.010	− 0.026	0.205	0.003	0.036

the interaction between price and promotion. Then we have the effect of the dummy variable that distinguishes between 64 and 128 oz pack sizes of detergents. This variable takes the value 1 for the 64 oz size. A holiday dummy variable is also included and its effect on category demand is presented next. This is followed by a set of parameters, σ, that correspond to the effects of store characteristics on the preferences of the brands in the category. Recall from our earlier discussion that these effects are the same for all the brands within the category. The coefficient 0.412 corresponding to the income variable implies that a higher income household has a greater probability of purchasing laundry detergents than a lower income household. The last set of parameters are the interactions between the store characteristics and

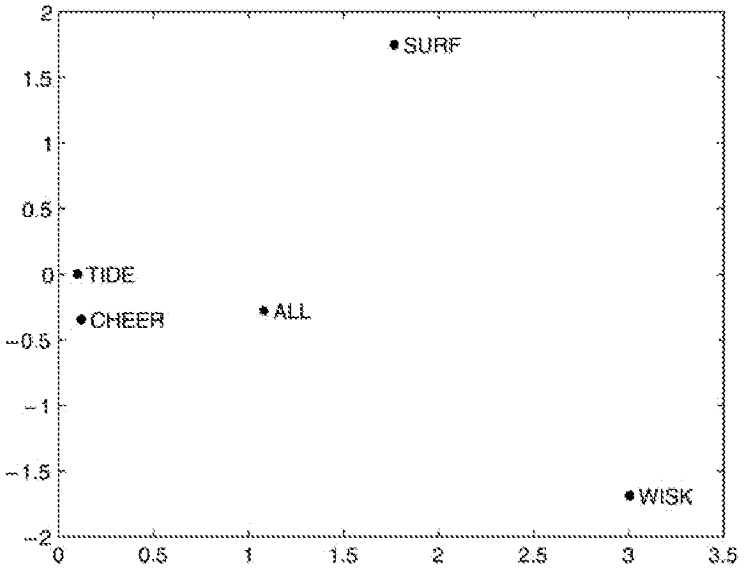

Figure 4. Perceptual map for laundry detergent.

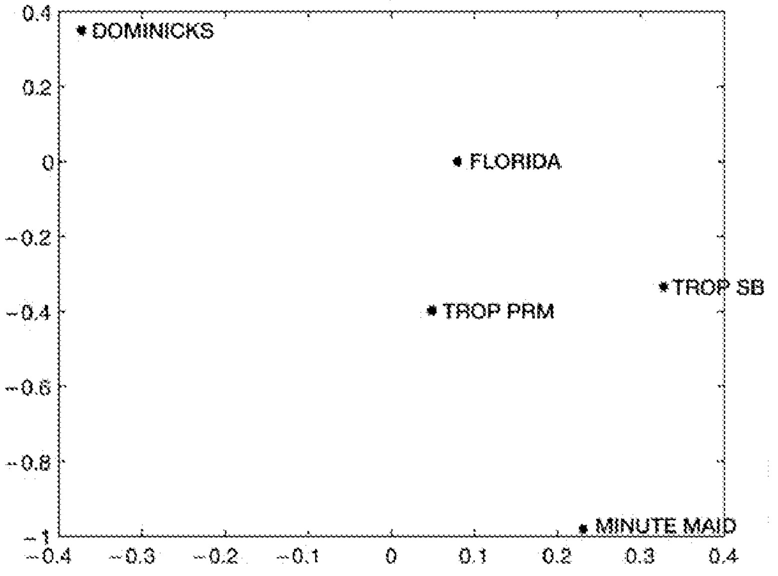

Figure 5. Perceptual map for refrigerated orange juice.

prices, γ. In other words, they represent the observed component of price heterogeneity across stores.[18]

In the laundry detergent category, table 4, we find that Tide is the highest-valued brand in the category. We also observe a preference for 64 oz versus 128 oz size packages. As expected, promotions increase the likelihood of purchase, as does the incidence of a holiday week. Price sensitivity has the correct negative sign and we do observe significant heterogeneity. Interestingly, we find the price sensitivity rises for promoted items. Using the estimated factors, from the covariance matrix of brand preferences, we plot the brand map for detergents in Figure 4. The map highlights the perceived similarities and differences across brands in the market. The horizontal axis seems to separate out brands manufactured by P&G (Tide and Cheer) from those manufactured by Unilever (Wisk, Surf and All). Further, the two P&G brands, Tide and Cheer are perceived to be similar to one another whereas the Lever brands appear a lot more differentiated in the minds of consumers. By positioning All close to the P&G brands, Lever may be using it as a "fighting" brand in the marketplace.

In the orange juice data, Table 5, we find that Tropicana premium has the highest mean preference effect, which is consistent with its dominant position in the category. We also observe a preference for the smaller 64-oz size, versus the larger 96 oz. As before, promotions increase the likelihood of purchase, but at the same time, they increase consumers' sensitivity to price. In Figure 5, we present the perceptual map for orange juice brands. Once again, we find the map of the brand preferences to be quite revealing. There appear to be three distinct groups of brands in the market. The first set, consisting of the Tropicana and Florida brands, is perceived as being different from the other brands. We do however, observe slight differences between the product not from concentrate (premium) and the product from concentrate (SB). The second group consists of Minute Maid, a Coca-Cola product, and the third group consists only of the store brand (Dominicks). This finding is good news for the national brands as it does appear that they have effectively differentiated themselves from the store brand.

In Table 6, we report the elasticities corresponding to the impact of the store characteristics on category size (the probability of purchase). Since store characteristics are allowed to interact with both preferences and price sensitivity, elasticity is a more revealing summary statistic. Median income seems to have a strong negative effect on category size for laundry detergents. In contrast, higher income areas are more likely to purchase refrigerated orange juice. Interestingly, the proportion of retired households tend to increase the size of both categories. Ethnicity has an almost negligible effect on both categories. The ability to shop increases the size of the laundry category; but, it has a very small negative impact on the size of orange juice. Both categories tend to be larger in areas with higher average

18 A no-purchase nested logit model generated a nest parameter near zero in each category. This finding is not surprising since we already allow for flexible substitution to no-purchase both across consumers (random brand intercepts) and across stores (interact mean intercept with store characteristics).

household values. Finally, neither of the competitive variables seem to have a strong effect on category size. Even doubling the distance from a competitor (a 100% increase in distance) will not generate more than about half a percent increase in the probability of purchase. This evidence suggests that the size of the inside share is driven primarily by demographic factors reflecting willingness-to-pay, rather than competitive or consumer search-related factors. Thus, we expect price discrimination to be the primary motive for zone pricing.

5.1. *Goodness-of-fit*

Our next objective is to try and assess the ability of our demand system to predict the observed store-level margins. An important distinction from previous work (e.g., Berry et al., 1995; Besanko et al., 1998; Sudhir, 2001b) is that we do not construct moment conditions based on the pricing model. Thus, our estimates do not rely on the supply-side. As a result, we can now test whether our supply-side, category pricing model, is appropriate. Using the observed wholesale prices, we compute the shelf prices for chain-level, zone-level and store-level pricing by solving the system of equations defined by (2) and (7). Since our main objective is to study the implications of pricing, we need to verify that our demand estimates and our category pricing model produce realistic measures of the retail category manager's ability to price above the wholesale price level.

The first step involves determining which model seems to come the closest to approximating the observed margins in the data. One way to think of this problem would be to solve a minimum distance procedure in which one minimizes the distance between true and estimated margins, using the covariance of the observed margins as a metric:

$$\min_{\mu}(\text{margin} - \mu)\Phi(\text{margin} - \mu), \tag{8}$$

where μ is the estimated margin and Φ is the covariance matrix of the observed margins. In Table 7, we take the margins implied by the store-level, zone-level and chain-level pricing policies and compute the corresponding criterion (8). The zone-pricing model seems to provide the best fit according to the minimum distance criterion. This result confirms our use of the reported zone configuration.

We now compare the mean predicted margins for each brand under the three pricing policies considered and compare these to the true margins observed in the data. To illustrate, in Tables 8 and 9, we report the Lerner index for each brand, $L = (\text{price} - \text{cost})/\text{price}$ which is the proportion of the shelf price attributed to the mark-up. These tables illustrate how closely each of the pricing strategies approximates the actual pricing behavior across the stores. While the median mark-up for each product is quite similar across the three pricing strategies, the important point is that they are fairly close to the actual mark-ups. Note that despite

Table 7. Minimum distance criterion.

	Laundry detergent	Refrigerated orange juice
Store	1.844	2.499
Zone	1.210	1.463
Chain	1.329	3.025

Table 8. Predicted margins for laundry detergent.

Brand	Size (oz)	Conditional share	TRUE	Store	Zone	Chain
Surf	64	6.2%	25.8%	19.5%	22.1%	22.3%
Wisk	128	7.0%	17.9%	22.8%	26.4%	26.6%
Wisk	64	14.1%	14.9%	21.5%	25.2%	25.4%
All	64	12.8%	22.3%	22.1%	24.2%	24.3%
All	128	10.9%	23.2%	23.9%	26.1%	26.2%
Cheer	64	6.3%	13.6%	14.4%	15.6%	15.7%
Cheer	128	5.3%	16.4%	15.1%	16.3%	16.4%
Tide	128	18.9%	14.7%	14.4%	15.5%	15.6%
Tide	64	18.4%	13.5%	13.8%	14.6%	14.8%

Table 9. Predicted margins for refrigerated orange juice.

Brand	Size (oz)	Conditional share	TRUE	Store	Zone	Chain
MM	64 oz	21.2%	23.7%	32.2%	30.6%	30.5%
MM	96 oz	4.1%	26.2%	30.1%	26.9%	26.5%
Dom	64 oz	24.4%	28.2%	38.7%	38.8%	38.4%
Trop Prm	64 oz	20.5%	28.1%	30.8%	28.4%	27.8%
Trop SB	64 oz	17.7%	30.1%	33.4%	30.7%	30.6%
Trop Prm	96 oz	7.7%	27.0%	28.9%	25.4%	24.7%
Florida	64 oz	4.3%	31.6%	41.8%	42.4%	42.1%

similarities in the means, the three models do not predict the same variance in prices by construction. Although not reported, the zone strategy tends to correlate the most highly with the observed data (similar to the minimum distance criterion above). For the laundry detergent data, the levels are reasonably close to the true values and the correlations are fairly high, especially for the largest-share items (Tide and All). For juice, our mean margins look, for the most part, fairly reasonable. The main exception is the store brand for which we over-predict the margin and for which price estimates are very poor. Our sense is that store-brand pricing may not fit well with the category management model. In fact, store-brands may well be priced below the category manager's level if the objective is to generate interest in the retailer's store brands in other product categories (Chintagunta, 2002). Nevertheless, the results confirm our use of the model to reflect the type of pricing behavior facing the managers of the chain.

6. Impact of micro-marketing pricing

Using the demand systems estimated in the previous section, we now set-out to measure the implications of the micro-marketing policies. First we assess the impact of the current zone-pricing scheme reported by Dominicks. Then, we use the database to construct a finer store-specific pricing strategy. Our analysis measures the impact on consumer welfare and retail profitability. Measuring consumer welfare allows us to assess the dollar value of losses or gains to consumers from the various pricing policies. Similarly, retail profits provide a dollar value of the losses and gains to the category manager.

6.1. Welfare implications for retailers and consumers

In Table 10, we report the total annual chain-wide welfare implications of various pricing policies for both of the categories studied. Recall that consumer welfare is measured as the Hicksian compensating variation, in equation (4) and that one interpretation of the computed consumer welfare number is the total dollar amount the chain would need to pay consumers to make them as well off under a new pricing policy as they were under uniform chain-wide prices. As expected, the move to store-level pricing generates substantial additional profits. While consumer welfare is reduced in refrigerated juice, it increases slightly in laundry detergent. Overall, profit

Table 10. Welfare implications of pricing policies.

Category	Aggregation	Profit	Change in profit	Change in profit ($)	Change in consumer welfare ($)	Total change in welfare ($)
Laundry Detergent	Chain	$1,148,500				
	Zone	$1,155,400	0.6%	$6900	$2158	$9058
	Store	$1,258,200	9.6%	$109,700	$16,215	$125,915
	Constrained store	$1,212,500	5.6%	$64,000	$39,381	$103,381
	Cluster	$1,192,500	3.8%	$44,000	$1,082	$45,082
Refrigerated orange juice	Chain	$3,336,000				
	Zone	$3,388,400	1.6%	$52,400	– $19,791	$32,609
	Store	$3,878,200	16.3%	$542,200	– $158,100	$384,100
	Constrained store	$3,582,400	7.4%	$246,400	$48,613	$295,013
	Cluster	$3,623,400	8.6%	$287,400	– $156,900	$130,500

gains exceed any consumer welfare losses so that aggregate chain-wide welfare increases.

For laundry detergent, we observe fairly small profit implications (only 0.6% gain) in going from chain to zone-pricing; although consumers do gain overall by $2158. Given the high-necessity oriented nature of laundry detergent, it seems intuitive that such an item would not exhibit tremendous variation in willingness-to-pay across store markets. In contrast, the orange juice category benefits reasonably well, generating a $52,400 (1.6%) gain in profits. At the same time, total consumer welfare falls $19,791. Recall from our discussion of the demand estimates in Table 6 that price discrimination appears to be the primary motivation for zone pricing and that competitive variables do not seem to generate much cross-store variation in either the slope or intercepts of demand. Accordingly we find that, on average, the price-differential between the chain and zone-level prices is most sensitive to income level—one of the key drivers of consumers' willingness to pay.[19]

Now we demonstrate how our demand system can be used to implement a more profitable store-specific pricing scheme. As above, we compare the prices and welfare levels from the zone model with those of the store-level model. As expected, the store-pricing leads to much larger profit gains than the zone pricing. We observe an almost $542,000 (16.3%) gain in profits relative to chain pricing for orange juice, and $109,700 (9.6%) for laundry detergent. At the same time, we observe $158,100 in losses to consumers in juice, which is small relative to the profit gains, but much larger than the losses from zone pricing. For detergent, consumers do in fact gain overall by $16,215. As before, we can see that the low consumer welfare losses are misleading once we look at each store. In Figure 6, we plot the compensating variation for orange juice (subplot A) and laundry detergent (subplot B) by store. Interestingly, we see that the negative chainwide impact on consumer welfare in juice is misleading since, in several stores, consumer surplus increases (recall that a negative valued compensating differential should be interpreted as consumer willing to pay to keep the new policy active—i.e., welfare-increasing). Similarly, in laundry detergent, we find the impact on consumers varies across stores; although the welfare changes are considerably smaller in this category. Similar to zone pricing, we find that the price differential between store and chain prices is most sensitive to income level and, to a much lesser extent, to the shopping index.[20] This similarity might suggest that zone pricing serves as a crude attempt to price discriminate by clustering stores according to similarities in consumer income levels.

The intuition for why consumers in some stores gain value while those in other stores do not relates to the ability of the store to re-align its product line pricing according to local demand. For instance, in store 75, most laundry detergent prices

19 Although not reported, we re-compute the prices after increasing each demographic variable by 10%. We find that a 10% increase in income level generates a 2% change in the differential.

20 We find that a 10% increase in income level generates a 2% change in the differential. We also find a 10% increase in the ability to shop has a 1% increase in the differential.

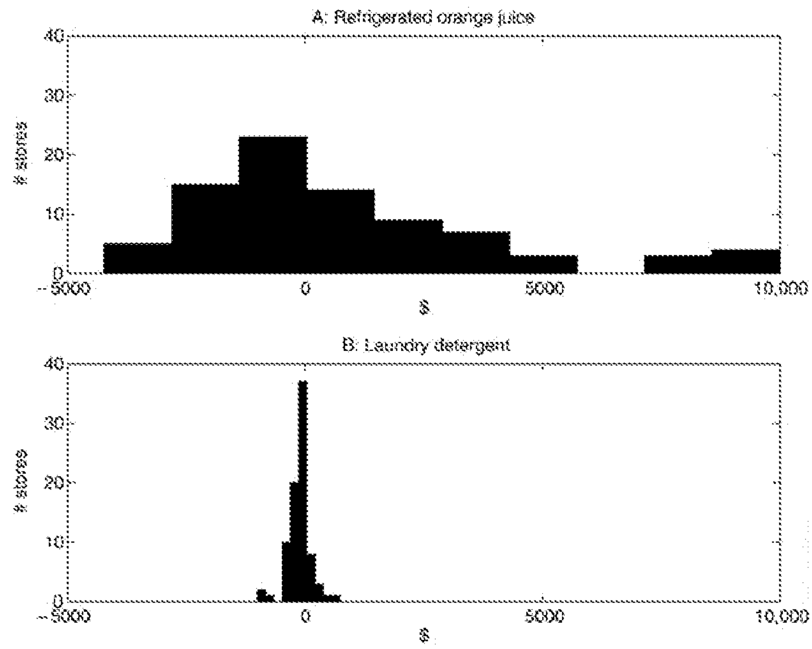

Figure 6. Welfare gains/losses from store-level pricing (Hicksian compensating differentials across stores).

are raised, on average, about 2–3% (e.g., 128 oz price rises by 26 cents). At the same time, the price of 64 oz Tide (the category leader in the store with roughly 15% more share than the 2nd-ranked alternative) is lowered by 1%. As a result, the conditional share of 64 oz Tide rises almost 16%, while the conditional shares of most of the other brands fall by 1–3%. Note that store 75 in zone 7, one of the high-price zones in which consumers have fairly high search costs and there are no nearby EDLP stores. As a result, demand is fairly inelastic, which explains why the store would want to raise its price level. Effectively, the store transfers surplus to consumers by making the highest quality brand slightly cheaper, while extracting surplus by raising the other products' prices without damaging their shares. In contrast, in store 128, 64 oz Tide has a much smaller lead, dominating the 2nd-ranked alternative by only 4% in category share. Interestingly, store 128 holds the prices of its top three products almost fixed, while lowering the prices of the remaining brands in the category 1–2%. As a result, shares become much more equalized in store 128, with the two largest-share goods falling 2–3%. At the same time, the conditional share of 64 oz Wisk rises almost 9%, making it (by a narrow margin) the new category leader, on average. Note that store 128 is in zone 11, which caters to households with relatively low incomes and house values. The area also caters to a higher proportion of ethnic households with larger families. As a result, demand is more elastic, which explains why the store reduces most prices, allowing for a better-value brand to gain more relative share.

An interesting question for customer relationship management involves which types of consumers end up better versus worse off after the chain adopts a more flexible pricing policy. In general, we find that both store and, to a much lesser extent, zone pricing decrease welfare in lower income neighborhoods. Similarly, welfare rises in areas with greater mean household values for all categories. We also observe welfare falling in areas with larger ethnic populations. In general, we do not find much relationship between age and welfare changes. Interestingly, zone-pricing appears to lower welfare in areas where households have a higher ability to shop. But, store-pricing raises welfare in stores catering to consumers with greater ability to shop. This outcome is not surprising since the store-pricing will clearly favor stores with more price-elastic demand. In terms of competition, the welfare implications of proximity to a Jewel varies across categories and are quite small. The impact of proximity to an EDLP store is even smaller. Overall, it would appear that demographics are more influential for driving patterns in welfare changes than proximity to competitors.

While our results showing greater profits from store pricing are consistent with those from Montgomery (1997), as noted before our analysis is different in nature. Most importantly, we note the trade-offs of zone and store-pricing in terms of the impact on consumer welfare. Since pricing decisions tend to be made by category managers, the analysis does not consider the impact of micro-marketing across all categories. The category manager sets prices assuming holding the valuation of the outside good (expenditures on other categories in the store) fixed. However, if all category managers adopt micro-marketing pricing policies, as consumer losses across categories accumulate, one might eventually expect losses in traffic as shoppers switch to other stores. We view store-switching as a longer-run implication based on store-wide prices and therefore beyond the scope of our analysis. Naturally, higher prices could also encourage new entry by a competing chain not currently in the territory. Lower prices, on the other hand, could evoke competitive retaliation, although they could also deter entry into the market. These issues are more long-term considerations that need to be balanced carefully with our short run findings if they are to be used as inputs into chain strategy. Based on our current results, it would appear that managers would be better off allowing for price discrimination in the laundry detergent category than refrigerated juice since the former category generates surplus both for the stores and the consumers.

6.2. Alternative recommended pricing policies

We now propose two alternative pricing policies. The first, addresses the issue of consumer welfare by balancing profits and consumer surplus. In particular, we compute the optimal prices and profits in each store after constraining consumers to obtain at least as much surplus as under a uniform chain-wide policy. The second pricing policy addresses possible computational and implementational concerns by store managers. We use the store-level results to construct a new five-zone configuration that still captures most of the benefits of the store-level pricing.

Store managers may be concerned that extracting too much value from consumers could generate store-switching in the long-run. Previous research has experimented with *ad hoc* constraints on pricing, such as holding the average price level fixed (Montgomery, 1997). The demand system we use generates a natural theory-based constraint—consumer welfare. We propose profit-maximizing store-level prices that are constrained to offer consumers at least as much surplus as a chain-level pricing policy. The prices under this policy will, by construction, make consumers at least as well off as under chain-pricing. Formally, the problem involves solving the following problem for each store s in each week t:

$$\max_{\{p_{sjt}\}_{j=1}^{J}} \Pi_{st} = \sum_{j=1}^{J} (p_{sjt} - w_{jt})Q_{sjt},$$

subject to the constraint:

$$\Delta W \geq 0,$$

where ΔW is computed as in equation (5).

Referring back to Table 10, we report the resulting change in profits and consumer welfare associated with such a policy in row labeled "constrained store". As expected, the constraint prevents the category manager from generating the same additional profits as under the unconstrained store-specific pricing policy of the previous section. However, even with the constraint, the manager is able to generate roughly half the gains of the store-pricing, an improvement in profits of 5.6% over a uniform chain-pricing policy, in the laundry detergent category, and 7.4% in the refrigerated orange juice. At the same time, overall consumer welfare rises, especially in the refrigerated juice category where unconstrained pricing led to overall losses to consumers.

A second consideration regarding the unconstrained pricing of the previous section is the potential complexity of coordinating a store-specific pricing policy across the 83 stores. Clearly, one advantage of a zone policy is the simplification of price-determination. The previous section demonstrates the ease with which an aggregate database can be leveraged to learn about differences in consumer willingness-to-pay. However, changing item prices in 83 stores could be costly from an implementation point of view. Therefore, we propose an alternative zone pricing policy. Using the store prices computed in the previous section, we construct share-weighted price indexes for both the refrigerated juice category and the laundry detergent category. Using the 83 price indexes, we then run a simple non-hierarchical cluster analysis[21] to generate five zones. Constraining prices to be the same across all stores within each of these five zones, we then re-compute the profits and welfare

21 We use the non-hierarchical k-means-based cluster function in Stata version 7.

levels that would prevail. In the row labeled "cluster" in Table 10, we find that this simple five-zone structure still generates substantial profit gains relative to the 16-zone pricing policy used by Dominicks during the time the data were collected. An interesting point is the fact that while the clusters offer notable gains to the retailer (relative to chain-pricing and the actual zone-pricing), consumers are better-off with 16-zone pricing policy used by Dominicks. Of course, the welfare-constrained policy proposed above could serve as a means of offsetting these losses to consumers. Relatedly, we find that this zone configuration is driven primarily by household values and, to a lesser extent, by household income.

7. Conclusions

Using a detailed database including weekly store-specific margins, we estimate flexible demand systems capable of generating reasonable approximations of the true data-generating process in two different product categories. We are able to rule out alternative explanations for the observed price dispersion across stores and conclude that the current zone-pricing is primarily a means of price discriminating based on geographic differences in consumer characteristics. We then use the demand system to simulate the benefits of switching from the current zone-pricing scheme to a data-based store-level pricing scheme. Interestingly, for a necessity item like laundry detergent, we find conservative gains in profits with small effects on consumers. However, for categories like refrigerated orange juice, which exhibits far more demand heterogeneity across stores, we find fairly large profit implications. At the same time, consumers experience differential welfare effects. In particular, we find that DFF's existing zone-pricing seems to target high prices to less affluent areas. Allowing DFF to use store-pricing exacerbates this effect. Interestingly, the shift to store pricing would also raise prices in areas where consumers are less able to shop. Finally, we illustrate how the firm could still generate profit gains even after constraining its prices to offer consumers a baseline level of surplus. Thus, we find that the chain could profitably implement store-level price discrimination without hurting its consumer population.

The estimation and validation of the demand system also raise a number of methodological issues. The parsimony of the logit demand system helps us avoid incorrect signs in the estimated cross-elasticity matrix, while still allowing for flexible substitution patterns. The instrumental variables procedure helps us resolve the endogeneity of prices due to omitted product attributes at the individual store level. Such endogeneity could, if ignored, lead to a downward bias in the estimated mean price sensitivity of demand. Since our data conta in the true store-level margins, we are then able to assess the proposed model in terms of its ability to reflect the chain's pricing conditions.

Our results add to the growing literature measuring the sources and welfare implications of price discrimination (e.g., Shepard, 1991; Leslie, 2001; Cohen, 2000, 2001; Iyer and Seetheraman, 2001). Unlike previous research, our data permit us to control more accurately for alternative explanations of price variation. Our access to margins also allow us to assess the predictive validity of our model and its reflection of pricing conditions.

In our current specification, we have only focused on the impact on consumers and retailers. An interesting area for future research is the impact on retail pass-through. While we expect wholesale prices to be determined at the market level, it may be of interest to see how zone pricing alters the extent that exogenous changes in wholesale prices are passed-through to consumers. Recently, Besanko et al. (2001b) document that, in the same data, retail pass-through rates are fairly high, especially for large-share items. They also find that pass-through rates vary across stores based on similar store characteristics as we use in our model. Using our structural model, we could measure the impact of zone-pricing on pass-through relative to a uniform chain-level price mechanism. This analysis could be interesting for manufacturers who are interested in identifying which consumers benefit the most from promotional wholesale discounts. Another area for future research is the impact of shifting to a micro-marketing pricing strategy on overall retail competition. While we rule out competition using proxies for local competitors, new datasets containing competitors' prices would allow for a more thorough investigation.

Acknowledgments

The data for this study were made available by Peter Rossi. The authors are very grateful to the editor for comments and suggestions. The authors also benefited from the comments of participants at the 2001 Choice Symposium hosted by U.C. Berkeley, the Marketing Science Conference at Wiesbaden, the Erasmus Research Institute in Management Conference, the INRA-IDEI Conference in Toulouse and workshops at Carnegie Mellon, Columbia, MIT, UCLA, Universidade Nova de Lisboa and Stanford Universities for comments and feedback. We also thank Tim Conley, Xavier Drèze, Alan Montgomery and K. Sudhir for comments and suggestions. The first two authors would also like to acknowledge the Kilts Center for Marketing at the University of Chicago for providing research funds. The second author acknowledges research support from the Beatrice Foods Faculty Fund at the Graduate School of Business. All correspondence may be addressed to the authors at the University of Chicago, Graduate School of Business, 1101 East 58th Street, Chicago, IL 60637 or via e-mail at jdube@gsb.uchicago.edu or pradeep.chintagunta@gsb.uchicago.edu.

References

Anderson, E. and N. Vilcassim. (2001). Retailer Category Pricing. Working Paper, University of Chicago.

Berry, S. (1994). "Estimating Discrete-Choice Models of Product Differentiation", Rand Journal of Economics 25, 242–262.

Berry, S., J. Levinsohn, and A. Pakes. (1995). "Automobile Prices in Market Equilibrium", Econometrica 63, 841–890.

Berto Villas-Boas, Sofia. (2001). Vertical Contracts Between Manufacturers and Retailers: An Empirical Analysis. Working paper, University of California, Berkeley.

Besanko, D., S. Gupta, and D. Jain. (1998). "Logit Demand Estimation Under Competitive Pricing Behavior: An Equilibrium Framework", Management Science (November).

Besanko, D., J.-P. Dubé, and S. Gupta. (2001a). Heterogeneity and Target Marketing Using Aggregate Retail Data: A Structural Approach. Working paper, Kellogg Graduate School of Management. Forthcoming, Management Science.

Besanko, D., J.-P. Dubé, and S. Gupta. (2001b). Identifying Retail Pass-through Rates for Competing Brands. Working paper, Kellogg Graduate School of Management. Forthcoming, Marketing Science.

Borenstein, S. (1991). "Selling Costs and Switching Costs: Explaining Retail Gasoline Margins", The Rand Journal of Economics 22, 354–369.

Chevalier, J.A., A.K. Kashyap, and P.E. Rossi. (2000). Why Don't Prices Rise During Peak Periods of Demand? Evidence from Scanner Data. Working paper, University of Chicago.

Chiang, J. (1991). "The Simultaneous Approach to the Whether, What, and How much to Buy Questions", Marketing Science 10, 297–315.

Chiang, R. and C. Spatt. (1982). "Imperfect Price Discrimination and Welfare", Review of Economic Studies 49, 153–181.

Chintagunta, P. (1993). "Investigating Purchase Incidence, Brand Choice and Purchase Quantity Decisions of Households", Marketing Science 12, 184–208.

Chintagunta, P. (1994). "Heterogeneous Logit Model Implications for Brand Positioning", Journal of Marketing Research 31, 304–311.

Chintagunta, P. (2002). "Investigating Category Profit Maximization at a Retail Chain", Journal of Marketing Research 39(2), 141–154.

Chintagunta, P., J.P. Dubé, and V. Singh. (2002). "Market Structure Across Stores: An Application of a Random Coefficients Model with Store Level Data". In Philip Hans Franses and Alan Montgomery (eds.), Advances in Econometrics: Econometric Models in Marketing. New York: JAI Press, 191–222.

Cohen, A. (2000). Package Size and Discrimination in Paper Towels. Working Paper, University of Virginia.

Cohen, A. (2001). Evidence of Price Discrimination Based on Package Size in the Paper Towel Industry. Working Paper, University of Virginia.

Dreze, X., S.J. Hoch, and M.E. Purk. (1994a). "Shelf Management and Space Elasticity", Journal of Retailing 70, 301–326.

Dreze, X., S.J. Hoch, and M.E. Purk. (1994b). "EDLP, Hi-Lo, and Margin Arithmetic", Journal of Marketing 58, 16–29.

Dubé, J.P. (2002). "Multiple Discreteness and Product Differentiation: Strategy and Demand for Carbonated Soft Drinks". Forthcoming, Marketing Science.

Elrod, T. (1988). "Choice Map: Inferring a Product Market Map from Panel Data", Marketing Science 7, 21–40.

Elrod, T. and M.P. Keane. (1995). "A Factor-Analytic Probit Model for Representing the Market Structure in Panel Data", Journal of Marketing Research 32, 1–16.

Hoch, S.J., B.-D. Kim, A.L. Montgomery, and P.E. Rossi. (1995). "Determinants of Store-Level Price Elasticity", Journal of Marketing Research 32, 17–30.

Iyer, G. and P.B. Seetheraman. (2003). To Price Discriminate or Not: The Choice of Station Type in Retail Gasoline Markets. Quantitative Marketing & Economics, forthcoming.

Kadiyali, V., P.K. Chintagunta, and N.J. Vilcassim. (2000). "Manufacturer-Retailer Channel Interactions and Implications for Channel Power: An Empirical Investigation of Pricing in a Local Market", Marketing Science 19, 127–147.

Kim, J., G.M. Allenby, and P.E. Rossi. (2002). "Modeling Consumer Demand for Variety", Marketing Science 21, 229–250.

Lal, R. and Carmen Matutes. (1994). "Retail Pricing and Advertising Strategies", Journal of Business 67, 345–370.

Leslie, P.J. (2001). Price Discrimination in Broadway Theatre. Working paper, UCLA.

McFadden, D. (1989). "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration", Econometrica 57, 995–1026.

McFadden, D. and K. Train. (2000). "Mixed MNL Models for Discrete Response", Journal of Applied Economerics 15, 447–470.

Montgomery, A.L. (1997). "Creating Micro-Marketing Pricing Strategies Using Supermarket Scanner Data", Marketing Science 16, 315–338.

Montgomery, A.L. and E.T. Bradlow. (1999). "Why Analyst Overconfidence about the Functional Form of Demand Can Lead to Overpricing", Marketing Science 18, 569–583.

Nevo, A. (2001). "Measuring Market Power in the Ready-To-Eat Cereal Industry", Econometrica 69, 307–340.

Nevo, A. (2000). "Mergers and Differentiated Products: The Case of the Ready-to-Eat Cereal Industry", The Rand Journal of Economics 31, 395–421.

Pakes, A. and D. Pollard. (1989). "Simulation and the Asymptotics of Optimization Estimators", Econometrica 57, 1027–1057.

Petrin, A. (2001). Quantifying the Benefits of New Products: The Case of the Minivan. Working paper, University of Chicago.

Reibstein, D.J. and H. Gatignon. (1984). "Optimal Product Line Pricing: The Influence of Elasticities and Cross-Elasticities", Journal of Marketing Research 21, 259–267.

Schmalansee, R. "Output and Welfare Implications of Monopolistic Third-Degree Price Discrimination", The American Economic Review 71, 242–247.

Shepard, A. (1991). "Price Discrimination and Retail Configuration", Journal of Political Economy 99, 30–53.

Singh, V. (2002). Competing With Mass Discounters: Implications for Retailer Pricing. Working paper, Kellogg Graduate School of Management.

Slade, M.E. (1995). "Product Market Rivalry with Multiple Strategic Weapons: An Analysis of Price and Advertising Competition", Journal of Economics and Management Strategy 4, 445–476.

Small, K.A. and H.S. Rosen. (1981). "Applied Welfare Economics with Discrete Choice Models", Econometrica 49, 105–130.

Sudhir, K. (2001a). "Competitive Pricing Behavior in the Auto Market: A Structural Analysis", Marketing Science 20, 42–60.

Sudhir, K. (2001b). "Structural Analysis of Competitive Pricing in the Presence of a Strategic Retailer", Marketing Science 20, 244–264.

Trajtenberg, M. (1989). "The Welfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners", Journal of Political Economy 97, 444–479.

Varian, H. (1985). "Price Discrimination and Social Welfare", The American Economic Review 75, 870–875.

Villas-Boas, M. and R. Winer. (1999). "Endogeneity in Brand Choice Models", Management Science 45(10), 1324–1338.

Villas-Boas, M. and Y. Zhao. (2001). The Ketchup Marketplace: Retailer, Manufacturers, and Individual Consumers. Working paper, UC Berkeley.

Zhao, Y. (2000). Price Dispersion in the Grocery Market. Working paper, University of Hong Kong.