

Category Pricing with State-Dependent Utility

Jean-Pierre Dubé, Günter J. Hitsch, Peter E. Rossi, Maria Ana Vitorino

Graduate School of Business, University of Chicago, Chicago, Illinois 60637
{jdube@chicagogsb.edu, gunty@chicagogsb.edu, peter.rossi@chicagogsb.edu, mvitorin@chicagogsb.edu}

There is substantial literature documenting the presence of state-dependent utility with packaged goods data. Typically, a form of brand loyalty is detected whereby there is a higher probability of purchasing the same brand as has been purchased in the recent past. The economic significance of the measured loyalty remains an open question. We consider the category pricing problem and demonstrate that the presence of loyalty materially affects optimal pricing. The prices of higher quality products decline relative to those of lower quality when loyalty is introduced into the model. Given the well-known problems with the confounding of state dependence and consumer heterogeneity, loyalty must be measured in a model which allows for an unknown and possibly highly nonnormal distribution of heterogeneity. We implement a highly flexible model of heterogeneity using multivariate mixtures of normals in a hierarchical choice model. We use an Euler equations approach to the solution of the dynamic pricing problem which allows us to consider a very large number of consumer types.

Key words: dynamic pricing; loyalty; state dependence; consumer heterogeneity

History: This paper was received May 9, 2006, and was with the authors 2 months for 3 revisions; processed by Kannan Srinivasan. Published online in *Articles in Advance* April 7, 2008.

Introduction

A large literature in marketing documents the presence of a particular form of state-dependent utility in choice among packaged goods items. The typical finding is that the utility of a brand is enhanced¹ if this brand was recently purchased, c.f. Erdem (1996), Roy et al. (1996), Seetharaman et al. (1999), Seetharaman (2004), and Horsky et al. (2006). Accordingly, consumers appear to exhibit loyalty to the previous brand purchased. While a great deal of work has been devoted to investigating various utility specifications and distributions of heterogeneity, comparatively little attention has focused on the implications of state dependence for substantive marketing decisions such as pricing.

We consider the implications of loyalty to the last brand purchased for the pricing problem facing a retailer managing a typical product category consisting of substitutable packaged good items. That is, we use the pricing problem and category profits as a metric for assessing the economic significance of the detected level of loyalty. We do not test whether loyalty influences observed prices per se. Rather, we compute the profit-maximizing prices that would arise from demand estimates calibrated from consumer panel data. Loyalty introduces a nontrivial source of dynamics into the category pricing problem. Current pricing decisions affect current product choices and, subsequently, the number of loyal

customers for each product in the following period. Thus, current prices influence future product demand. For the static category pricing problem, prices are typically high for products with relatively high perceived qualities and, hence, high consumer willingness to pay. The solution to the dynamic pricing problem requires the category manager to control the flow of customers to those products where the future returns to loyalty are highest. If the returns to loyalty are highest for relatively high quality products, then state dependence can introduce a counteracting downward pressure on the prices of higher quality goods.

We formulate the dynamic pricing problem facing the retailer as a dynamic program in which the retailer chooses a Markovian price policy function, which provides the optimal price at time t as a function of the state of the system at time t . Aggregate demand at time t is the sum of the demands over various consumer types and states of loyalty to various brands. The “state” of the system is a vector of the proportions of consumers of each type and each loyalty state. For any realistic problem, this is a very high dimensional state vector which renders many of the standard value function approximation methods (c.f. Judd 1998, Chapter 12) impractical. To solve the problem numerically, we exploit the fact that there is a stationary long-run pricing policy and we solve for the optimal prices using the Euler equations which this stationary pricing policy must satisfy. This insight enables us to use Euler equations to solve a dynamic programming problem with a high dimensional state space.

¹ This form of state dependence is typically termed “brand loyalty.”

The infinite horizon pricing problem is appropriate for situations in which firms expect to remain in business for more than a few periods. In finite horizon problems, as the terminal period approaches, the incentives for the firm to invest in loyal customers decrease and make the pricing solutions nonstationary. This means that these models, taken literally, would predict price trends of the sort which we typically do not see in retailer price data. In the theoretical literature on brand loyalty (c.f. Villas-Boas 2004), it is common to consider a two-period problem. While this can be a useful simplification to illustrate some of the economic forces at work in these problems, it is not appropriate for application to retail pricing problems. The only empirical work on dynamic pricing with brand loyalty (Che et al. 2005) considers only finite horizon problems as well. Our methodological contribution consists of an approach for solving single-agent infinite horizon problems which, thus far, proved too difficult to solve for realistic empirical models of demand.

While our approach correctly solves the single-agent dynamic pricing problem, it does not extend to competitive environments like the one considered by Che et al. (2005). Accommodating competition increases the computational complexity considerably and, in most instances, renders the solution intractable. Rather than solve the competitive dynamic pricing problem, Che et al. (2005) approximate the solution by imposing a simplification of the state space to reduce the dimensionality of their problem.² However, it is unknown whether this approach can approximate the solution to the true dynamic pricing problem satisfactorily. In addition, Che et al. (2005) consider a finite horizon problem which does not allow for the computation of long-run or stationary pricing policies.

In order to claim that our optimal pricing results have relevance for marketing practice, we require a realistic utility specification as well as a flexible distribution of consumer heterogeneity. Fortunately, we can draw on the large literature for advice on the specification of the utility model, as authors have tried and tested a wide variety of specifications. In the specification of the distribution of heterogeneity, the current state of the art (c.f. Keane 1997 and Seetharaman et al. 1999) is a normal distribution (an exception is Allenby et al. 1998). The degree of measured loyalty has been noted to be highly sensitive to the manner in which heterogeneity is modeled. Since there is little theoretical or empirical guidance as to the

nature of the distribution of heterogeneity, we adopt a semiparametric approach in which the distribution of utility coefficients is approximated by a mixture of multivariate normals.

We fit our demand model to scanner data from the margarine and refrigerated orange juice categories. Computation of optimal prices taking into account brand loyalty reveals considerable differences between forward and non-forward-looking firm behavior. For our estimated demand values, accounting for loyalty results in higher firm profits. More generally, if the firm maximizes profits over the infinite horizon, the prices of different quality goods will change relative to the one-period optimal prices. In particular, the price gap between brands of low and high perceived quality will be reduced. This is driven by the fact that the marginal profitability of a customer loyal to the high quality good is higher than for a low quality good. Hence, a forward-looking firm uses prices to control the number of loyals to the higher quality good.

Model

Demand

We consider a market where households exhibit state dependence in product choice. Product demand is derived by aggregating over individual household decisions. Households choose among J products and an outside option. The utility index from product j at time period t is:

$$u_{jt} = \alpha_j + \eta p_{jt} + \gamma I\{s_t = j\} + \varepsilon_{jt}, \quad (1)$$

where p_{jt} is the product price and ε_{jt} is the standard iid error term used in most choice models. In the model given by Equation (1), the brand intercepts represent vertical product differentiation. The error terms ensure that there will be a degree of “horizontal” differentiation between products in the sense that the aggregate market share will be positive for all products. $s_t \in \{1, \dots, J\}$ summarizes the history of past purchases from the perspective of impact on current utility. We take the standard approach, often termed the “state-dependent” model of demand, and assume that the history can be characterized by an indication of which brand was purchased on the last purchase occasion (e.g., Erdem 1996 and Seetharaman et al. 1999). If $\gamma > 0$, it can be interpreted as a household’s “loyalty” to the specific brand. If a household buys product k in period t , then $s_{t+1} = k$, i.e., the household becomes loyal to that brand. If the household chooses the outside option, then $s_{t+1} = s_t$, i.e., the household’s brand loyalty remains unchanged.

All demand parameters $\Theta = (\alpha_1, \dots, \alpha_j, \eta, \gamma)$ can be household specific. There are N household types in

² The approximation uses aggregate market shares as a proxy for the state vector. The correct state vector contains the fraction of consumers of each type loyal to each brand. This reduces the size of the state vector from $N \times J$ to J (if there are N types of consumers and J brands) but has unknown approximation properties.

the market. The behavior of each household type n is fully characterized by the taste vector Θ_n . We assume that for each household type, there is a continuum of consumers with mass 1 in the market. This assumption is for convenience because it makes the demand system and the transition of the entire state vector deterministic from the perspective of the firm.

The market is summarized by a state vector x_t which is comprised of subvectors of dimension J for each of the N consumer types: $x_t = (x_t^1, \dots, x_t^N)$ and $x_t^n = (x_{1t}^n, \dots, x_{jt}^n)$. x_{jt}^n is the fraction of consumers of type n who are loyal to product j at time t . Since consumers must be loyal to one of the products, $\sum_j x_{jt}^n = 1$. Let $\Pr_j^n(s_t^n, p_t) = \Pr_j(s_t^n, p_t | \Theta_n)$ denote the choice probability for product j given the state and the vector of prices p_t . Demand for product j and household type n is given by:

$$D_j^n = \sum_{k=1}^J x_{kt}^n \Pr_j^n(s_t^n = k, p_t).$$

Aggregate demand for product j is given by summing over types using the distribution of households over types and loyalty states:

$$D_j(x_t, p_t) = \sum_{n=1}^N D_j^n = \sum_{n=1}^N \sum_{k=1}^J x_{kt}^n \Pr_j^n(k, p_t). \quad (2)$$

Evolution of the State Vector

If households exhibit brand loyalty, the pricing problem becomes more complicated because current period pricing decisions influence the loyalty state and thus affect the future state of the market. In particular, the model given by Equation (1) implies that the states evolve according to a Markov chain. The last period's purchase defines the loyalty state from this point on, giving rise to the Markovian property. The transition matrix of the Markov chain is populated by the conditional choice probabilities. For example, if a household is currently loyal to product k , then the household will remain in state k next period if it purchases product k . Thus, the conditional choice probabilities are the state transition probabilities. We must remember, however, that if the household purchases the outside good, then its loyalty state remains unchanged. Consequently, we must add the conditional probability of choosing the outside good to the diagonal elements of the transition matrix. Let Q denote the transition matrix where element $q_{k,j}$ represents the probability of transitioning from state j to state k :

$$Q_{k,j}^n(p_t) = \begin{cases} \Pr_k^n(j, p_t) + \Pr_0^n(j, p_t) & \text{if } k = j \\ \Pr_k^n(j, p_t) & \text{if } k \neq j. \end{cases} \quad (3)$$

The transition matrix is a function of the current price vector and is indexed by household type. The state

vector for household types evolves according to the standard Markov chain equation:

$$x_{t+1}^n = Q^n(p_t)x_t^n. \quad (4)$$

Since standard choice models deliver strictly positive choice probabilities, the elements of Q are positive, guaranteeing that the chain given by Equation (4) has a unique stationary distribution for a fixed price vector. For ease of reference to the Markov decision problem literature, we can also write Equation (4) as:

$$x_{t+1} = f(x_t, p_t). \quad (5)$$

Dynamic Category Pricing

The pricing problem is posed as an infinite horizon problem in which the retailer maximizes the discounted flow of profits subject to the constraint imposed by the state evolution equation:

$$V(x_0) = \max_{\{p_t\}} \sum_{t=0}^{\infty} \beta^t \omega(p_t, x_t) \\ \text{s.t. } x_{t+1} = f(x_t, p_t) \quad t = 0, \dots, \infty. \quad (6)$$

$\omega(p_t, x_t)$ is the current period profit given by:

$$\omega(p_t, x_t) = \sum_{j=1}^J (p_{jt} - c_j) D_j(x_t, p_t). \quad (7)$$

β is the discount factor (<1) and c_j is the marginal cost of product j (assumed constant over time for simplicity).

We are solving the infinite horizon version of the retailer's problem in Equation (6). Some might argue that it is more realistic to solve a finite period problem as most retailers are not literally infinitely lived. There are two problems with this approach. First, if we pick a finite number of periods, we must pick a very large number of periods to approximate any real retail application since most retailers expect to stay in business for many years. There is no simple way to solve, for example, a 40-period problem except via a "brute force" nonlinear programming approach which would require optimizing a constrained nonlinear program of very high dimension. Second is the problem of terminal effects. It is well-known the optimal solution reflects the approach of "cap T" in a finite horizon model. In our case, this would distort the incentives to control the brands to which customers are loyal. The infinite horizon problem provides a good approximation to the behavior of retailers in a stationary environment where there are no approaching terminal conditions.

To solve Equation (6), we must pick an optimal infinite sequence of price vectors or, equivalently, find a value function V that solves the Bellman equation:

$$V(x) = \max_p \{\omega(p, x) + \beta V(f(x, p))\}. \quad (8)$$

Given a solution to the Bellman equation, we can find the policy function which yields optimal prices as a function of the state vector $p_t = \sigma(x_t)$. The optimal prices maximize the right-hand side of the Bellman equation.

To solve the functional Equation (8), the standard approach involves approximating the value function V via some interpolation method and using a value or policy function algorithm (Judd 1998). We are solving for a policy function which maps the $N \times J$ dimensional state space into a J dimensional vector of prices. For even a small number of products and types, this problem quickly becomes computationally intractable. We wish to apply these methods to situations in which there are potentially hundreds of types of consumers and five or more products. For this reason, we must exploit the special structure of our problem in order to render the computation of optimal prices feasible.

Steady State Prices via Euler Equations

We focus on computing a steady state with constant prices and a reproducing state vector. As we show below, a steady state pricing policy exists for our problem. Even though the steady state pricing policy does not theoretically guarantee a constant level of steady state prices, we can nevertheless establish existence of the latter by construction, i.e., we solve the Euler equations for a steady state level of price. In Appendix A, we provide details on the assumptions necessary for a steady state to exist and for our approach to work.

We can solve for the steady state combination of prices and state vector using a Euler equation approach applied to the Bellman equation (see, for instance, Miranda and Fackler 2002). In steady state, we can write the Bellman equation as:

$$V(x) = \omega(x, p) + \beta V(f(x, p)), \quad (9)$$

where $p = \sigma(x)$ and $x = f(x, p)$. Differentiating Equation (9) with respect to both x and p and applying the envelope theorem yields a set of three equations which the steady state must satisfy:

$$\begin{aligned} \frac{\partial \omega}{\partial p} + \beta \left[\frac{\partial f}{\partial p} \right]^t \lambda &= 0 \\ \frac{\partial \omega}{\partial x} + \beta \left[\frac{\partial f}{\partial x} \right]^t \lambda &= \lambda \\ x &= f(x, p). \end{aligned} \quad (10)$$

Here, λ is an NJ vector of derivatives of the value function V with respect to the state vector. Thus, a steady state is the three-tuple (x^*, p^*, λ^*) which solves Equation (10). Given the very high dimension of λ , it would be useful if Equation (10) could be simplified

so that we only have to find roots of a smaller number of equations.

If $(I - \beta[\partial f/\partial x]^t)$ is invertible, we can solve for λ and substitute this out of the system in Equation (10):

$$\frac{\partial \omega}{\partial p} + \beta \left[\frac{\partial f}{\partial p} \right]^t \left(I - \beta \left[\frac{\partial f}{\partial x} \right]^t \right)^{-1} \frac{\partial \omega}{\partial x} = 0 \quad (11)$$

$$x = f(x, p). \quad (12)$$

Equation (11) is an equation in vector notation which represents J roots. Equation (12) is simply the steady state condition that the state vector must reproduce itself. This suggests the following algorithm for finding the steady state solution:

1. Pick x_0 .
2. Solve for p_0 given x_0 using Equation (11).
3. Solve for x_1 given p_0 using Equation (12).
4. Repeat until $\|p_k - p_{k-1}\| < tol$.

In our problem, $\partial f/\partial x = Q$. So $(I - \beta Q^t)$ must be invertible to pursue this approach. Since Q^t is the transition matrix for a Markov chain, its eigenvalues are all ≤ 1 . Since $\beta < 1$, it is not possible for $(I - \beta Q^t)x = 0$ unless x is 0. The invertibility condition is, therefore, satisfied. In addition, we can compute the solutions to Equation (12) given p by simply computing the largest eigenvector of Q . This provides the basis for a very fast algorithm to compute the steady state. Appendix B provides the derivatives of the payoff function with respect to the price and state vectors.

A “brute” force solution to the Euler equations in Equation (10) would require finding the roots to $2NJ + J$ nonlinear equations. By substituting out for λ and exploiting the Markov chain structure, we have reduced the problem to a manageable dimension (only J roots). It would be impractical to solve Equation (10) with $N > 100$, especially many thousands of times as we do in our application below.

Econometric Specification

We specify a multinomial logit model for each household type and set the intercept and price of the outside good to zero. The probability that household h chooses alternative j given loyalty to product k is given by:

$$\begin{aligned} \Pr(j | s = k, \theta^h) \\ = \frac{\exp(\alpha_j^h + \eta^h \text{Price}_j + \gamma^h I\{j = k\})}{1 + \sum_{l=1}^J \exp(\alpha_l^h + \eta^h \text{Price}_l + \gamma^h I\{l = k\})}. \end{aligned} \quad (13)$$

To accommodate differences across households, we use a potentially large number of household types and a continuum of households of each type. Our approach will be to specify a very flexible but continuous model of heterogeneity and then use draws from the posterior of this model as “representative”

of the large number of consumer types. Each household in our data will be viewed as “representative” of a type. We will use Markov chain Monte Carlo methods to construct a Bayes estimate of each household’s coefficient vector.

We use a mixture of normals as the distribution of heterogeneity in a hierarchical Bayesian model.³ With enough components in the mixture, we will be able to accommodate deviations from normality as well as to make our conclusions resistant to outliers. Let θ^h be the vector of choice model parameters (see Equation (1)) for household h . The mixture of normals model specifies the distribution of θ^h across households as follows:

$$\begin{aligned}\theta^h &\sim N(\mu_{ind}, \Sigma_{ind}) \\ ind &\sim \text{multinomial}(\pi),\end{aligned}\tag{14}$$

where π is a vector giving the mixture probabilities for each of the k components. We implement posterior inference for the mixture of normals model of heterogeneity and the multinomial logit base model as in Rossi et al. (2005). We note that it would be a simple matter to add observables such as demographics to Equation (14) as in Rossi et al. (2005). In our empirical application, we did not find that demographics contribute much to explaining differences between consumers.

Our MCMC algorithm will provide draws of the mixture probabilities as well as the normal component parameters. Thus, each MCMC draw of the mixture parameters provides a draw of the entire multivariate density of household parameters. We can average these densities to provide a Bayes estimate of the household parameter density. We can also construct Bayesian credibility regions for any given density ordinate to gauge the level of uncertainty in the estimation of the household distribution.

Some might argue that you do not have a truly non-parametric method unless you can claim that your procedure consistently recovers the true density of parameters in the population of all possible households. In the mixture of normals model, this requires that the number of mixture components (K) increases with the sample size. Our approach is to fit models with successively larger numbers of components and gauge the adequacy of the number of components by examining the fitted density as well as the Bayes factor associated with each number of components. What is important to note is that our improved MCMC algorithm is capable of fitting models with a

large number of components at relatively low computational cost.

The mixture of normals is a direct generalization of the normal models typically assumed in the literature. Another method commonly used is a finite mixture model. Neither the mixture of normals nor the finite mixture models enforce a restriction that the price coefficient be negative, which would be necessary for finite profits and optimal pricing. In typical applications, this is not a source of concern as the estimates of price coefficients are usually negative. However, a more satisfactory approach might be to employ a prior which only puts mass on negative values. This can be achieved either by using truncated normals (as in Boatwright et al. 1999) or via reparameterization. We leave this for future research.

Empirical Results

Demand Estimation

We now present the results from the estimation of the mixture-of-normals random coefficients logit using the refrigerated orange juice and the 16 oz. tub margarine shopping data. We use panel data collected by ACNielsen for 2,100 households in a large Midwestern Scantrack market between 1993 and 1995. The data contain all purchase behavior for the refrigerated orange juice and the 16 oz. tub margarine categories in the two largest supermarket chains. In each category, we focus only on those households that purchase a brand at least twice during our sample period. Table 1 lists the products considered in each category as well as the purchase incidence, product shares, and average retail/wholesale prices. Over 85% of the trips to the store recorded in our panel data do not involve purchases in the product category. This means that the outside good share is very large as is typical in many product categories and analyses of scanner data. In addition, households who adopt a pattern of purchasing a product on a regular cycle will be perceived as relatively price insensitive as the changes in price of the category relative to the outside good will have little influence on purchase incidence for these households.

In Table 2, we report the log marginal density for several alternative model specifications for each category. The posterior probability of a model specification is monotone in the log marginal density, so that by choosing the model with the largest log marginal density, we are picking the model with the highest posterior probability (c.f. Rossi et al. 2005). By comparing models with and without loyalty and with varying types of heterogeneity, we can assess the importance of incorporating loyalty and nonnormality. For example, in the margarine category, we observe that more than one normal component is

³ The combination of heterogeneity and state dependence raises a potential initial conditions problem. Consistent with past literature on consumer state dependence, we implicitly assume the initial state s_0 is exogenous and that the joint distribution of heterogeneity is independent of s_0 .

Table 1 Description of Data

Product	Retail price	Wholesale price	Percent of trips
Refrigerated orange juice category			
64 oz. MM	2.21	1.36	1.52
Premium 64 oz. MM	2.62	1.88	0.96
96 oz. MM	3.41	2.12	2.01
Premium 64 oz. TR	2.73	2.07	3.96
64 oz. TR	2.26	1.29	0.93
Premium 96 oz. TR	4.27	2.73	1.09
No purchase (percent of trips)		89.53	
No. of households		355	
No. of trips per household		89.51	
No. of purchases per household		9.37	
Margarine category			
Promise	1.69	1.22	2.87
Parkay	1.63	1.02	1.10
Shedd's	1.07	0.83	2.78
Blue Bonnett	1.42	0.76	0.33
ICBINB	1.55	1.11	5.13
No purchase (percent of trips)		87.80	
No. of households		455	
No. of trips per household		81.02	
No. of purchases per household		9.89	

Notes. MM—Minute Maid; TR—Tropicana; ICBINB—I Can't Believe Its Not Butter.

required to achieve a good fit. In addition, Table 2 confirms that consumer demand for frequently purchased consumer packaged goods (CPG) products exhibits state dependence even after controlling for heterogeneity with a very flexible and nonnormal distribution of tastes. In Appendix C, we find that richer specifications including demographics, promotional variables, or brand-specific loyalty coefficients do not lead to much improvement in model fit.

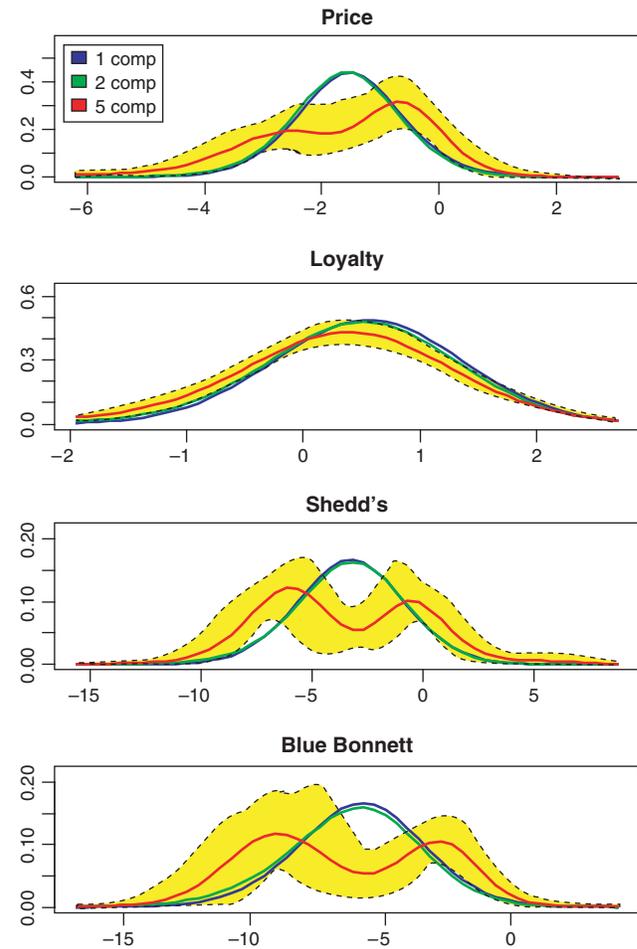
We now assess the extent of nonnormality of the fitted distributions of taste parameters. Ultimately, our goal is to estimate the distribution of tastes across households, not to attach any meaning or substantive significance to the parameters of the mixture components.

In Figure 1, we plot several fitted densities from the 1-, 2-, and 5-component mixture models for the

Table 2 In-Sample Fit and the Role of Heterogeneity and State Dependence

Model	No. of Components	Log marginal density	
		Margarine	OJ
No loyalty	No heterogeneity	-19,278.99	-15,888.78
	5-component	-13,780.92	-11,622.32
Loyalty	No heterogeneity	-16,524.03	-14,330.43
	1-component	-13,855.71	-11,503.13
	2-component	-13,827.89	-11,510.86
	5-component	-13,739.29	-11,505.56
	10-component	-13,727.12	-11,478.65

Figure 1 Fitted Densities for Margarine Price, Loyalty, and Brand Intercepts

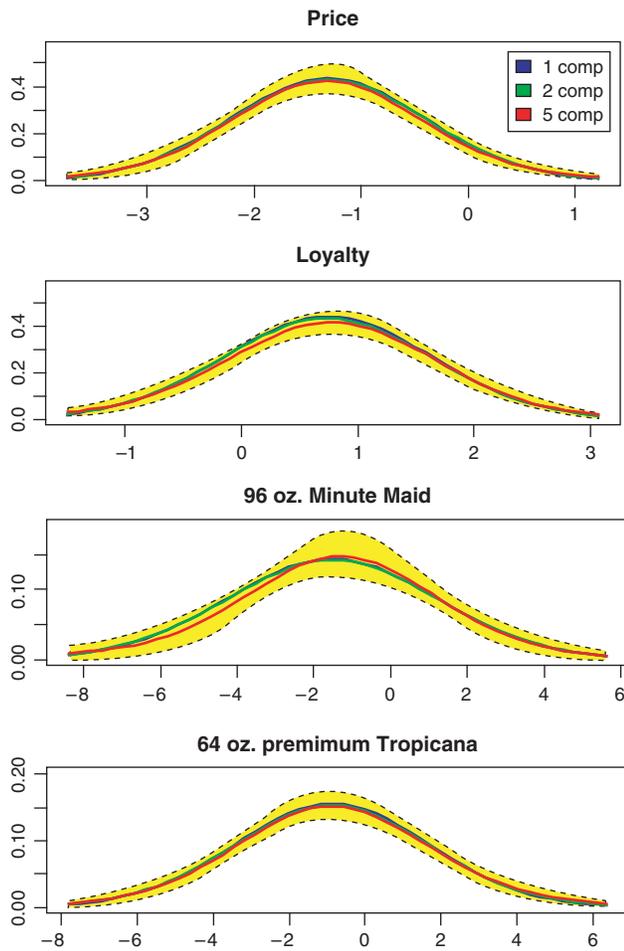


margarine data. We also report the 95% posterior credibility region for the 5-component mixture model. This region provides point-wise evidence for the nonnormality of the population marginal density for a given coefficient. For the price coefficient (upper panel), the 5-component model leads to a very asymmetric and slightly bimodal density. In contrast, a symmetric 1-component model has both a mode and tails lying outside the credibility region for the 5-component model.

Figure 1 provides compelling evidence of the need for a flexible model capable of accommodating nonnormality. In both the upper and lower panels, the brand intercepts from the 5-component model exhibit striking bimodality that simply can not be captured by the 1- or 2-component models. The bimodality exhibited in Figure 1 has the interpretation that there are households who differ markedly in their quality perceptions for margarines. In general, the results suggest that one would recover a very misleading description of the data-generating process if the usual symmetric normal (1-component) prior were used to fit these data.

Copyright: INFORMS holds copyright to this *Articles in Advance* version, which is made available to institutional subscribers. The file may not be posted on any other website, including the author's site. Please send any questions regarding this policy to permissions@informs.org.

Figure 2 Fitted Densities for Orange Juice Price, Loyalty, and Brand Intercepts



In Figure 2, we report analogous fitted densities from the orange juice category. These plots illustrate why we do not get the same improvement from more mixing components as we did in the margarine data. For the orange juice data, the 1-component normal approximation seems adequate. We note that Bayesian semiparametric procedures such as the one used here do not “overfit” the data in the sense of fitting slight deviations from the data. Even the 5-component mixture model returns estimated densities quite close to normality for the orange juice category.

Dynamic Category Pricing

We now use the estimated 5-component demand systems for orange juice and margarine to study the economic implications of state dependence (i.e., loyalty). Each sample household is treated as one of the N consumer types, where N is the total number of sample households used. A consumer type’s tastes are estimated by using the posterior mean of the households’ choice model coefficients. We trim those households with positive posterior mean price coefficients, which eliminates one household in the

Table 3 Tub Margarine (16 oz.) Category Pricing Results: Loyalty vs. No Loyalty

	16 oz. tub margarine		Brand intercepts (mean of posteriors)
	With loyalty beta = 0.999	No loyalty	
Promise	4.21	3.97	-2.62
Parkay	1.62	1.72	-5.01
Shedd’s	1.17	1.15	-3.65
Blue Bonnett	1.40	1.55	-6.47
ICBINB	2.59	4.20	-2.16
Outside share	0.92	0.94	
SS profit (\$)	40.18	35.49	

Note. SS—steady-state.

margarine category and 13 households in the orange juice category. Aggregate demand is obtained by summing over household “types” as in Equation (2).

We use category pricing as our metric of economic performance. For each category, we compute two sets of stationary category prices using the demand estimates with and without loyalty. In the latter, we imagine that the retailer calibrates a misspecified demand system with no loyalty and then uses this demand system to maximize category profits. Since there is no loyalty included in this demand system, prices are set to maximize the static category profits. We contrast the no loyalty case to the case where a retailer is aware of the presence of loyalty and properly calibrates a demand system with state dependence. In this case, prices are set to maximize the net present value of category profits in steady state with a discount factor of $\beta = 0.999$. We then compare category profits for these two vectors of prices using the demand system fit with loyalty. Since the demand specification with loyalty provides a superior fit and the highest posterior model probability, we refer to this as the “true” demand system.

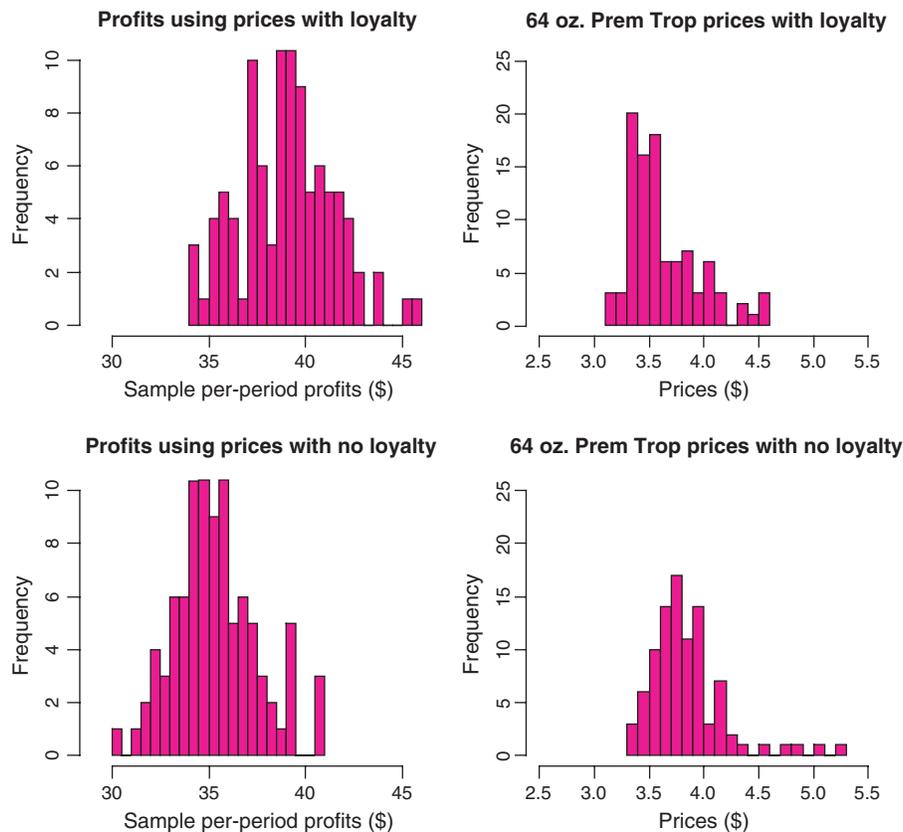
Tables 3 and 4 report the stationary price levels for each product using both demand with and without loyalty. The tables also report the corresponding sample level per-period category profits and the category

Table 4 Refrigerated Orange Juice Category Pricing Results: Loyalty vs. No Loyalty

	Refrigerated orange juice		Brand intercepts (mean of posteriors)
	With loyalty beta = 0.999	No loyalty	
MM 64	2.60	2.39	-2.33
MM 64 PR	3.23	2.97	-2.50
MM 96	3.76	3.18	-1.62
TR 64 PR	3.38	3.85	-0.84
TR 64	2.20	2.04	-2.62
TR 96 PR	4.38	4.52	-1.13
Outside share	0.88	0.92	
SS profit (\$)	33.17	31.69	

Copyright: INFORMS holds copyright to this *Articles in Advance* version, which is made available to institutional subscribers. The file may not be posted on any other website, including the author’s site. Please send any questions regarding this policy to permissions@informs.org.

Figure 3 Posterior Distribution of Prices and Category Profits: Refrigerated Orange Juice



outside share. These calculations were obtained by aggregating demand over each consumer type with parameters set to the posterior mean for that household. In the profit calculations, we measure costs as the mean wholesale list prices of each product using Leemis data on wholesale list prices. In both categories, profits are considerably lower when steady-state prices do not account for loyalty. In the orange juice category, the stationary per-period profits are 4% lower when loyalty is ignored. The profits are over 11% lower in the margarine category.

Our goal is to explore the implications of loyalty or state dependence for optimal pricing. The results reported in Tables 3 and 4 confirm that there are important differences in pricing policy that arise from the presence of loyalty. These computations were undertaken using the posterior mean of household parameters as the Bayes estimator for each household “type.” There is a good deal of parameter uncertainty in each of these household point estimates. To understand the implications of this uncertainty for optimal prices and profits, we use draws from the posterior distributions of household parameters. The posterior distribution of optimal profits is induced by the joint posterior distribution of household parameters:

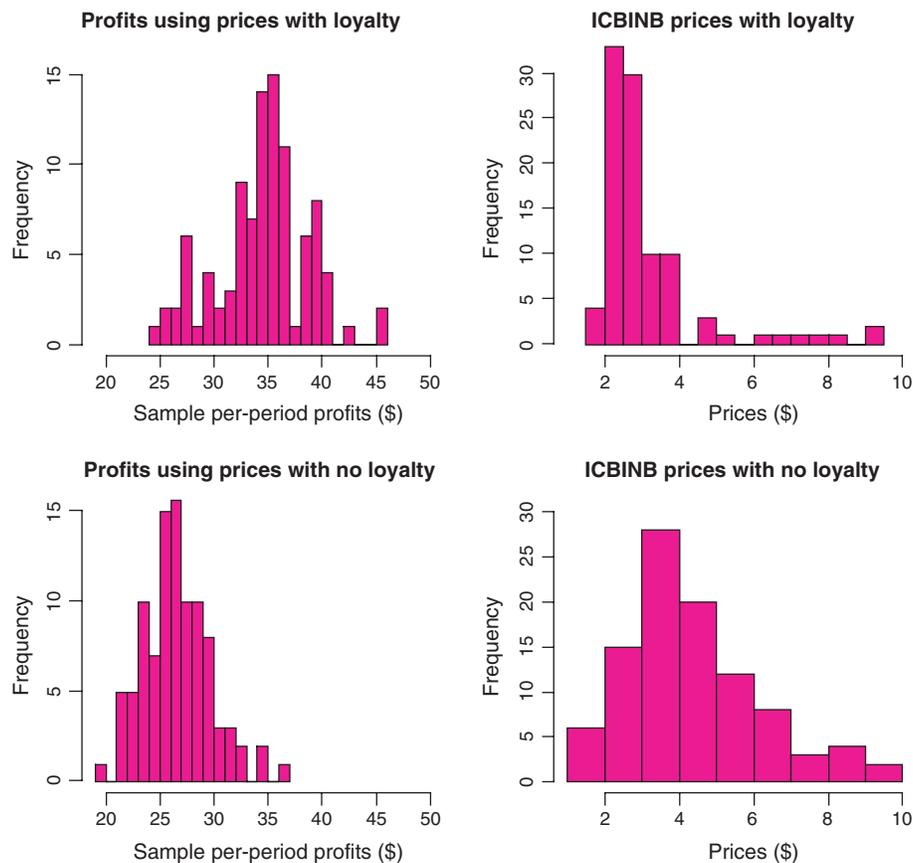
$$\Pi(p^* | \theta_1, \dots, \theta_H) = \sum_{t=0}^{\infty} \beta^t \omega(p^*(\theta_1, \dots, \theta_H), x^*).$$

Here p^* is the optimal stationary price which is a function of the household parameters and costs, and ω is the one period payoff function defined in Equation (8). Note that the stationary state vector is also an implicit function of the household parameters since this vector solves the equation $x^* = f(p^*, x^*)$. Given that we have draws for the household parameters, it is possible to compute the posterior distribution of aggregate profits.

In Figures 3 and 4, we report the full posterior distribution of profits⁴ in each category under the “with loyalty” and “with no loyalty” cases. In both categories, we see a decline in profits even after accounting for uncertainty around our demand parameters. Disregarding loyalty in long-run pricing decisions results in a net value loss to the category manager. We should note that even though individual households may have imprecisely estimated parameters,

⁴ We note that the support of the posterior distribution of the price coefficient includes positive values. There is a very small mass on this tail but nonetheless it exists. This means that the posterior distribution of profits has a small mass on an infinite level of profits. Comparison of two posterior distributions as in Figures 5 and 6 is still valid even with some support on infinite values (we could not, however, compute moments of these distributions). For optimal price computations, we delete draws with positive price coefficients.

Figure 4 Posterior Distribution of Prices and Category Profits: 16 oz. Tub Margarine



this does not mean that we estimate aggregate profitability imprecisely. Figures 3 and 4 show that with an economically meaningful profit metric, we can discern the impact of acknowledging loyalty even in the presence of considerable household level uncertainty.

The overall loss in value to each category is attributable to qualitative differences in how the category manager sets prices in the presence of loyalty. Tables 3 and 4 suggest that, after accounting for loyalty, the category manager lowers the prices on the highest average perceived-quality products relative to those of lower quality products. For example, “I Can’t Believe It’s Not Butter (ICBINB)” in the margarine category and “Premium 64 oz. Tropicana” in the orange juice category have the highest mean brand intercepts. When loyalty is not accounted for, the manager sets much higher prices for these items. Figures 3 and 4 confirm this finding by reporting the entire posterior distributions of these prices under the loyalty and no loyalty cases.

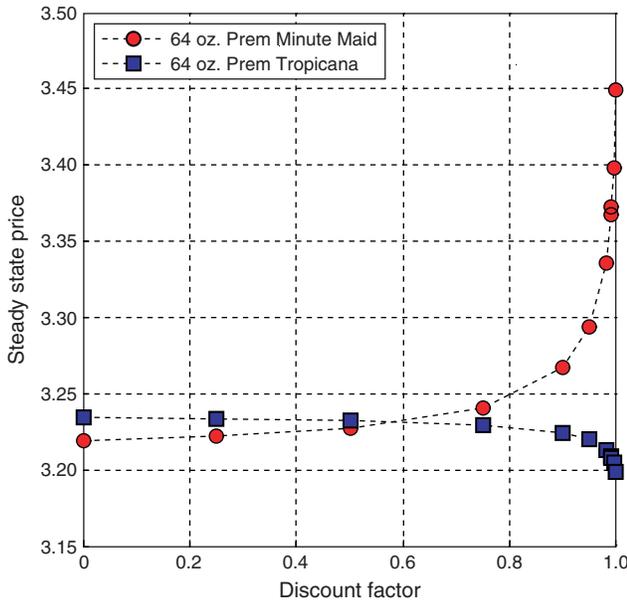
At first glance, these pricing patterns seem counter-intuitive. With loyalty, the dynamics of the category require a forward-looking manager to use prices not only to collect revenues, but also to control the state (i.e., the flow of loyal customers). The demand model

allows for both horizontal and vertical product differentiation. The nonlinearity of the demand model implies that a higher quality (i.e., vertically differentiated) item has a higher marginal profitability from loyalty. A forward-looking retailer will use prices to migrate customers from lower quality to higher quality brands in order to reap the greater future benefits from loyalty to the latter. Some consumers will continue to purchase the lower quality products based on their horizontal characteristics. A myopic retailer will set higher prices on the high quality brands as she does not anticipate the benefit of creating customers loyal to the high quality brands.

Exploring Different Parameter Settings

In this section, we explore the theoretical underpinnings of the pricing results discussed in the previous section. By varying either the discount factor or the level of the state-dependence coefficient, we can manipulate the magnitude of the long-run “future” profit implications of current prices. By varying the degree of vertical product differentiation, we can also manipulate the incentives to the retailer to use prices to control the flow of loyal customers. To simplify our discussion, we consider a model with only two products, 64 oz. premium Minute Maid and 64 oz. premium Tropicana, where the latter is found on average

Figure 5 The Effect of the Discount Factor on Relative Prices

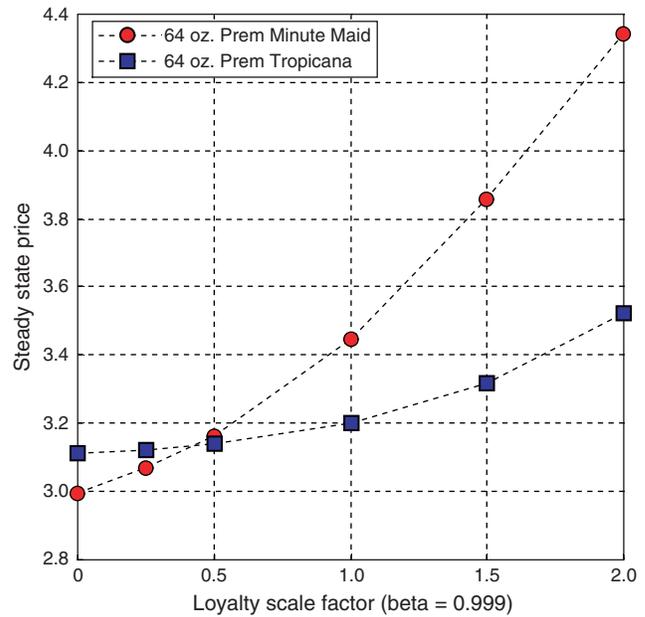


to have a higher perceived quality. For simplicity, we assume they both have the same costs (i.e., wholesale prices).

In Figure 5, stationary price levels are plotted against the discount factor. For a low discount factor, we find that the manager assigns a higher price to the product with the higher perceived quality, 64 oz. premium Tropicana. However, for relatively large discount factors, the price differential shrinks. In fact, for very high discount factors, the manager lowers the price of the high quality good below that of the low quality good.⁵ At these very high discount factors, the long-run effects of pricing have a bigger impact on long-run profits. The motive to use prices to control consumers' product loyalties puts downward pressure on the high quality good's price. Hence, a more patient retailer puts more weight on the long-run effect of prices and, hence, fundamentally changes the relative prices in the category to push consumers toward the higher quality product. In Appendix C, we show that these results are robust to richer empirical specifications including brand-specific loyalty coefficients and controls for features and displays.

In Figure 6, we vary the degree of loyalty. At each point along the x axis, we have multiplied each consumer type's estimated loyalty coefficient by the corresponding scale (which varies from 0 to 2). For low levels of loyalty, as we increase the degree of loyalty, we observe a shrinking of the price differential between the high and low quality goods. However, as loyalty becomes sufficiently large, we see once again

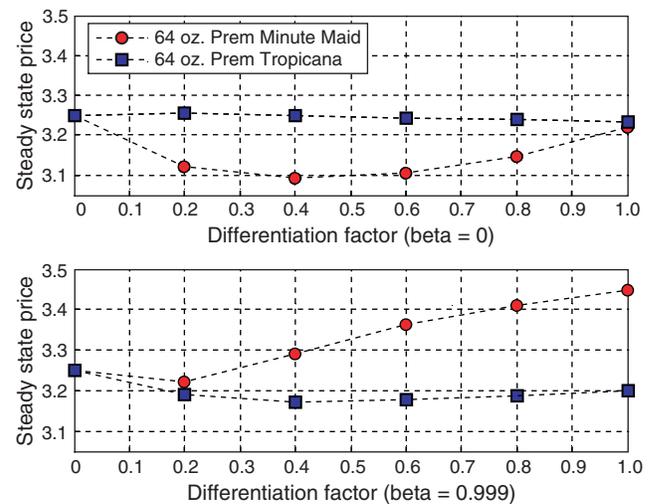
Figure 6 The Effect of Loyalty on Relative Prices



a switch in the rank order of prices. The intuition is the same as in the case of raising the discount factor. As loyalty becomes larger, the manager becomes more concerned about the long-run effects of prices.

In Figure 7, we vary the degree of vertical product differentiation by varying the difference in the brand intercepts. If the products did not exhibit any asymmetries, loyalty would not introduce any real dynamics into the model. Rather, both products would have equal amounts of loyal customers so that loyalty would simply amount to an outward shift in demand. At each point along the x axis in both panels, we have set the difference in each household's taste for products 1 and 2 equal to the actual estimated difference multiplied by the corresponding scale. In the upper

Figure 7 The Role of Product Differentiation on Relative Prices



⁵ The per-period profits are also higher than when $\beta = 0$ because the price of the low quality good is raised.

Copyright: INFORMS holds copyright to this *Articles in Advance* version, which is made available to institutional subscribers. The file may not be posted on any other website, including the author's site. Please send any questions regarding this policy to permissions@informs.org.

panel, we set the discount factor to zero and in the lower panel we set the discount factor to 0.999.

In both panels, when the demands are symmetric (i.e., no difference in the brand intercepts), we obtain identical prices for the myopic manager ($\beta = 0$) and for the forward-looking manager ($\beta = 0.999$). Hence, dynamics play no role once we have symmetric demands. However, as we increase the asymmetry (i.e., quality differential) between the two products, we see an increasing gap in the prices of the two goods. For the myopic manager, the high quality product has the relatively high price. However, for the forward-looking manager, the low quality product has the relatively high price.

Conclusions

Our results suggest that loyalty is both statistically discernable and has meaningful economic consequences. The estimated loyalty values have a large influence on category pricing. Ignoring loyalty leads to lower long-run total category profits. In addition, our findings suggest that a rational category manager should set category prices differently in the presence of loyalty. In particular, the manager may reduce the relative price of the high quality items. This change in pricing structure arises from the combination of consumer loyalty on the demand side and forward-looking category pricing on the supply side. At first glance, these relative pricing results may seem counterintuitive—lower margins on the higher quality products than in the static case without state-dependent demand. However, here we must recognize that the retailer seeks to maximize the net present value of total category profits. Customers who are loyal to the high quality (lower price elasticity) products are more valuable than customers loyal to the lower quality products. Thus, the investment motive (in loyals) causes the retailer with long-run profit maximization objectives to lower prices on higher quality brands relative to the case in which there is no loyalty effect.

By exploiting the Markovian structure of the model, we can compute stationary price levels directly from the Euler equations of the category profit maximization problem. We circumvent the usual curse of dimensionality that limits approaches requiring the evaluation of the optimal pricing policy. This approach enables us to consider a much richer scope of differentiated products and heterogeneous consumer types. As a result, we can base our pricing calculations on real demand estimates from scanner panel data.

Our approach of using Euler equations to compute the optimal steady state prices for dynamic models can be used to tackle other more complicated models of state dependence such as those that include multiple lags of the loyalty variable or that include a “wear out”

effect for brand loyalty. Thus, future research could probe more deeply into the sources of consumer brand loyalty and their implications for pricing. We believe that this is an exciting opportunity for future research. In a very general sense, we have made computations possible for a large class of retail dynamic pricing models which, heretofore, have eluded solution.

A limitation of the use of Euler equations is that they can not be used to solve dynamic models in competitive environments. The literature has yet to develop methods for solving competitive dynamic pricing problems in the context of rich differentiated products demand systems like the ones estimated herein. We view this as an important direction for future research.

Acknowledgments

The authors benefited from the comments of seminar participants at Northwestern University. They also thank Aviv Nevo for comments and suggestions. The first three authors are grateful to the Kilts Center for Marketing at the University of Chicago for financial support. Dubé acknowledges the Beatrice Foods Co. faculty fund for research support.

Appendix A

Existence of a Steady State and Solution to the Dynamic Pricing Problem

In this section, we prove the existence of an optimal pricing policy and provide a sufficient condition for the existence of a steady state with a corresponding price level $p^* = \sigma(x^*)$.

The state space for the category pricing problem, $X \subseteq \mathbb{R}^N$, is a convex and compact set. We restrict the firm's choice of prices to some bounded rectangle $p \in A = [0, a_1] \times \dots \times [0, a_j]$. Each a_j needs to be finite but can be chosen to be arbitrarily large. Therefore, the current period profit function $\omega(p, x)$ is bounded and continuous on $A \times X$. By Theorem 4.6 in Stokey and Lucas (1989), there is a unique value function $V: X \rightarrow \mathbb{R}$ that satisfies the Bellman Equation (8).

In order to prove the existence of a steady state, we make the following assumption:

ASSUMPTION. For each state x , there is a unique price that maximizes the right-hand side of the Bellman Equation (9).

Define the mapping $G(x) = Q(\sigma(x))x$. Assuming the optimal pricing policy function is used, G maps the current state x_t into next period's state $x_{t+1} = G(x_t)$. Applying the theorem of the maximum, σ is continuous and hence G is also continuous. Furthermore, G is a mapping from the convex, compact set X into itself. Brower's fixed point theorem then implies that G has a fixed point $x^* = G(x^*)$. This fixed point is a steady state of the dynamical system corresponding to our pricing problem.

Note that, in general, the steady state may not be unique. Also, given that we allow for a very flexible distribution of heterogeneity in tastes, the assumption that there is a unique optimal price at each state may be violated. In that case, a steady state need not exist. In practice, we find x^* and the corresponding optimal price $p^* = \sigma(x^*)$ using the algorithm proposed above. In our application, we were always able to compute a steady state and did not find any evidence for nonuniqueness.

Appendix B

Derivatives of Profit Function and Transition Probabilities

Equation (11) provides the Euler equation whose roots can be used to find the stationary price and state vector. In this appendix, we provide the details about the computation of the necessary derivatives in this equation.

We must first recognize that the single period payoff or profit function is the sum of the single period profit functions for each household type. In addition, the state vector evolves according to a block-diagonal Markov chain:

$$\omega(x_t, p_t) = \sum_{i=1}^N \omega^{(i)}(x_t, p_t)$$

$$x_{t+1} = \begin{bmatrix} x_{1,t+1} \\ \vdots \\ x_{N,t+1} \end{bmatrix} = f(x_t, p) = \begin{bmatrix} f^{(1)}(x_{1,t}, p) \\ \vdots \\ f^{(N)}(x_{N,t}, p) \end{bmatrix}$$

$$= \begin{bmatrix} Q^1(p_t)x_{1,t} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & Q^N(p_t)x_{N,t} \end{bmatrix}$$

Thus, the Euler equation can be written as the sum of various differentials for the N types of households:

$$\sum_{i=1}^N \left(\frac{\partial w^{(i)}}{\partial p} + \beta \frac{\partial f^{(i)}}{\partial p} (I - \beta Q_i)^{-1} \frac{\partial w^{(i)}}{\partial x} \right)$$

We will provide the formulas for each of these terms, suppressing the redundant type superscript:

$$\frac{\partial w}{\partial x} = \frac{\partial}{\partial x} (x^t (\Pr(p))^t (p - c)) = \Pr(p)^t (p - c)$$

Here, $\Pr(p)$ is the matrix of conditional choice probabilities:

$$\Pr(p) = [\Pr(i | \text{loyal to } j)] = \frac{\exp(x_{i,j}^t \theta)}{1 + \sum_l \exp(x_{l,j}^t \theta)}$$

$$\frac{\partial w}{\partial p} = \frac{\partial}{\partial p} ((p - c)^t \Pr(p)x) = \Pr(p)x + (p - c)^t \left[\frac{\partial}{\partial p} \Pr(p)x \right]$$

The derivative of $\Pr(p)x$ is a $J \times J$ matrix with elements given by:

$$\frac{\partial}{\partial p} (\Pr(p)x) = \left[\sum_{k=1}^J \frac{\partial}{\partial p_j} (\Pr(i | \text{loyal to } k)x_k) \right],$$

with

$$\frac{\partial \Pr(i | \text{loyal to } k)}{\partial p_j} = \begin{cases} -\Pr(i | k) \Pr(j | k) \eta \\ \Pr(i | k) (1 - \Pr(i | k)) \eta \end{cases}$$

Finally, $\partial f / \partial p$ (also a $J \times J$ matrix for each type) can be obtained in a similar manner:

$$\frac{\partial f}{\partial p} = \frac{\partial}{\partial p} (Q(p)x) = \left[\sum_{k=1}^J \frac{\partial}{\partial p_j} (q_{i,k}) x_k \right],$$

with

$$\frac{\partial}{\partial p_j} (q_{i,k}) = \begin{cases} \frac{\partial}{\partial p_j} (\Pr(i | \text{loyal to } k)) & \text{if } i \neq k \\ \frac{\partial}{\partial p_j} (\Pr(i | \text{loyal to } i) + \Pr(\text{outside good} | \text{loyal to } i)) & \text{if } i = k \end{cases}$$

and

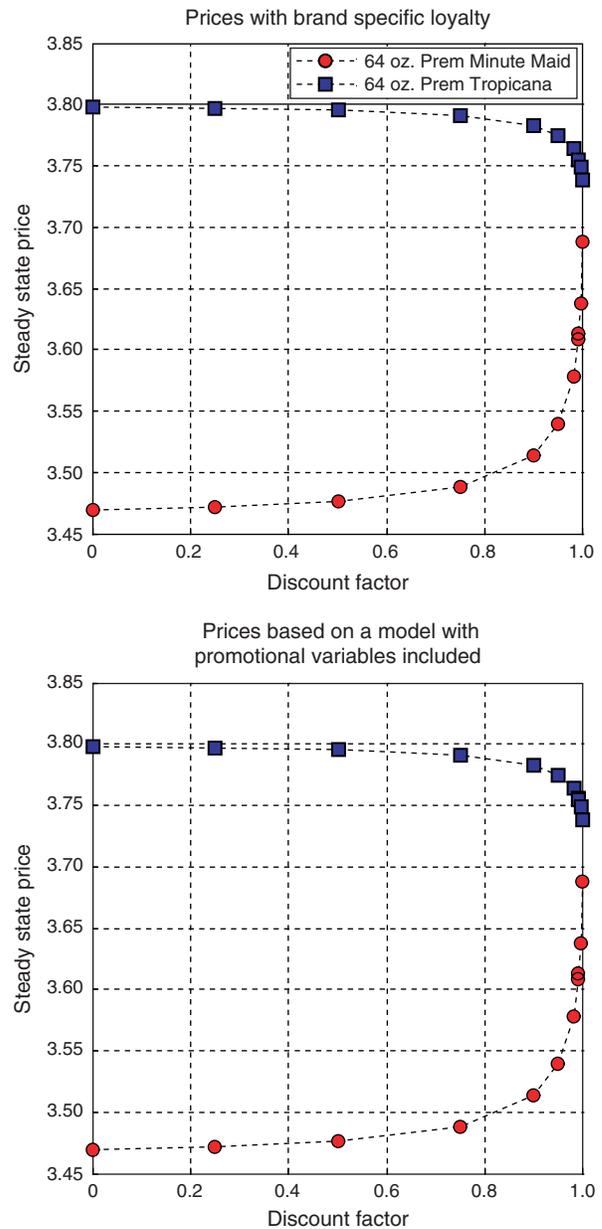
$$\frac{\partial}{\partial p_j} (\Pr(\text{outside good} | \text{loyal to } i)) = -\Pr(\text{outside good} | \text{loyal to } i) \Pr(j | \text{loyal to } i) \eta$$

Appendix C

Robustness Checks

Our demand model in Equation (13) is a simple specification that has been used widely in the empirical literature. This model includes price and loyalty terms and unobserved heterogeneity in all coefficients. It is certainly possible to complicate this model by the addition of interactions with demographic variables, promotional variables such as

Figure C.1 Robustness Checks



feature and display, and the use of product-specific loyalty coefficients (in the spirit of Guadagni and Little 1983). We have estimated each of these more elaborate models. The addition of standard demographic variables (household income, household size, and a dummy for female head age) add little to the fit and do not change the distribution of loyalty across households. Addition of brand-specific loyalty coefficients does make a modest (<1%) difference in the log marginal density or posterior model probability. Addition of feature and display variables also makes a modest (about 1%) improvement in fit. However, the real question is whether or not the solutions to the dynamic pricing problem are sensitive to these changes in model specification. In Figure C.1, we provide pricing results for the model with brand-specific loyalty (top panel) and the model with promotional variables (bottom panel). These results should be compared with our simple model results in Figure 5. The qualitative results remain unchanged. We still find that the relative price of high quality to low quality products is reduced.

References

- Allenby, G. M., N. Arora, J. L. Ginter. 1998. On the heterogeneity of demand. *J. Marketing Res.* **35** 384–389.
- Boatright, P., R. McCulloch, P. E. Rossi. 1999. Account-level modeling for trade promotion: An application of a constrained parameter hierarchical model. *J. Amer. Statist. Assoc.* **94** 1063–1073.
- Che, H., K. Sudhir, P. B. Seetharaman. 2005. Pricing behavior in markets with state dependence in demand. *J. Marketing Res.* Forthcoming.
- Erdem, T. 1996. A dynamic analysis of market structure based on panel data. *Marketing Sci.* **15** 359–378.
- Guadagni, P. M., J. D. C. Little. 1983. A logit model of brand choice calibrated on scanner data. *Marketing Sci.* **1** 203–238.
- Horsky, D., S. Misra, P. Nelson. 2006. Observed and unobserved preference heterogeneity in brand choice models. *Marketing Sci.* **25** 322–335.
- Judd, K. L. 1998. *Numerical Methods in Economics*. MIT Press, Cambridge, MA.
- Keane, M. P. 1997. Modeling heterogeneity and state dependence in consumer choice behavior. *J. Bus. Econom. Statist.* **15** 310–327.
- Miranda, M. J., P. L. Fackler. 2002. *Applied Computational Economics and Finance*. MIT Press, Cambridge, MA.
- Rossi, P. E., G. M. Allenby, R. McCulloch. 2005. *Bayesian Statistics and Marketing*. John Wiley and Sons, New York.
- Roy, R., P. K. Chintagunta, S. Haldar. 1996. A framework for investigating habits, “The Hand of the Past,” and heterogeneity in dynamic brand choice. *Marketing Sci.* **15**(3) 280–299.
- Seetharaman, P. B. 2004. Modeling multiple sources of state dependence in random utility models: A distributed lag approach. *Marketing Sci.* **23** 263–271.
- Seetharaman, P. B., A. Ainslie, P. K. Chintagunta. 1999. Investigating household state dependence effects across categories. *J. Marketing Res.* **36**(4) 488–500.
- Stokey, N. L., R. E. Lucas. 1989. *Recursive Methods in Economics*. Harvard University Press, Cambridge, MA.
- Villas-Boas, J. M. 2004. Consumer learning, brand loyalty, and competition. *Marketing Sci.* **23** 134–145.