Switching Costs and Store Choice

Timothy J. Richards and Jura Liaukonyte*

Selected Paper prepared for presentation at the
2018 Agricultural & Applied Economics Association Annual Meeting
Washington, D.C., August 5-August 7

May 3, 2018

Abstract

Switching costs are generally regarded as anti-competitive as firms can raise prices to "locked-in" consumers, at least up to the cost of switching to a lower-priced alternative. However, there is some evidence, both theoretical and empirical, that tends to show the opposite. Namely, suppliers, anticipating the pool of rents potentially available, compete aggressively to acquire non-switching consumers. Moreover, fixed shopping costs and uncertain prices imply that there is a "real option" value embedded in consumers' shopping behavior, and which must be compensated if consumers are to switch stores. We argue that retail prices are lower when retailers use programs designed to increase customer retention, or "stickiness." We test our theory using a panel of household-level store-choice data. Contrary to the conventional wisdom, we find that loyalty is pro-competitive and leads to lower prices than would otherwise be the case. We also find that approximately 50% of the loyalty effect is attributable to the existence of a real option in store choice.

Keywords: retail prices, food retailing, shopping-basket model, switching costs.
JEL Classification: D12, D83, L13, L81.

*Richards is the Marvin and June Morrison Chair of Agribusiness, in the Morrison School of Agribusiness, W. P. Carey School of Business, Arizona State University, Mesa, AZ; Liaukonyte is associate professor in the Charles H. Dyson School of Applied Economics and Management, Cornell University, Ithaca, NY. Contact author: Richards. email: trichards@asu.edu; ph. 480-727-1488. address: 7231 E. Sonoran Arroyo Mall, Mesa, AZ. 85212. Support from the Agricultural and Food Research Initiative, National Institute for Food and Agriculture (AFRI-NIFA), USDA is gratefully acknowledged.

1 Introduction

Store choice, because it represents a pattern that is repeated frequently out of need, and often driven by convenience, involves a high degree of habituation (Bell, Ho, and Tang 1998; Leszczyc, Sinha, and Timmermans 2000). Because of this observation, the cost of switching stores is substantial: Consumers have to find their preferred item in a new store, perhaps try store brands that are unique to the new store, and maybe join a new frequent-shopper program. While orthodox economic reasoning considers even small switching costs to be anti-competitive (Klemperer 1987; Farrel and Klemperer 2007), more recent work shows that the opposite can be true. Rhodes (2014a), Villas-Boas (2015) and Cabral (2016) show, and argue theoretically, how switching costs can potentially be pro-competitive, even in a market with frequently-purchased consumer goods. At the very least, stores have an incentive to compensate consumers for switching through lower prices (Arie and Grieco 2014). We argue that the compensation argument of Arie and Grieco (2014) is more general in that switching consumers must be compensated for not only the fixed costs of switching, but the real option that switching costs create. Because real option values tend to create an important gap between prices that would induce new consumers to try a store, and higher prices that would lead them to exit, we tend to observe much greater apparent loyalty than would be the case in a purely competitive market. We maintain that this phenomenon is responsible for the observation that food retailers – supermarkets – tend to be intensely competitive, even after waves of mergers have left few national chains standing. In this paper, we examine this hypothesis using a panel of store-choice data and a structural econometric model of retail competition.

Shoppers who are loyal to one store expect to pay lower prices for identical items in their home store relative to another. While this would seem to be obvious if the store is chosen for offering consistently low prices on preferred items, our observation is important as it is sustained during periods when the loyal store does not have the lowest prices. That is, when switching would seem to make sense, it does not occur, and the consumer pays higher prices.

However, over many purchase cycles, the price paid in the preferred store tends to be lower than that paid elsewhere. Although this observation is not conclusive as prices paid between stores may differ for a number of reasons, it is nonetheless important as it is inconsistent with the anti-competitive effects of switching costs.

Multi-product retailers exist because shopping involves fixed costs of search that can be efficiently allocated over a larger number of items purchased on each trip (Bettencourt and Gautschi 1990; Smith 2004; Rhodes 2014b). Once a consumer has committed to one store, therefore, the costs of re-optimizing the search process, and finding a new store can be substantial. Whether consumers use a sequential search process (Wildenbeest 2011) or a fixed-sample procedure (de Los Santos, Hortacsu, and Wildenbeest 2012), the optimal choice of a store still results from comparing the marginal costs of search, to the marginal benefit of receiving lower prices for the entire basket, or perhaps a preferred mix of items and service attributes. Based on empirical evidence of inertia in store choice (Rhee and Bell 2002), the expected net benefit of re-optimizing for each shopping trip is apparently not sufficient to make store-switching a general pattern of behavior, as store loyalty is more rather the norm. If switching costs in store-choice are large, therefore, then it seems reasonable to ask if they lead to above-normal profits for supermarkets?

There is an emerging body of theoretical literature that suggests this might not be the case. Retailers understand that their customers are forward-looking, and face substantial costs of switching loyalty to a different store. In this setting, retailers face a "harvest or invest" pricing decision – harvest rents from loyal customers in the short run through higher prices, or invest in building the cohort of loyal customers in the long run by reducing prices. If most customers are loyal, as was once assumed, then the harvest effect dominates and prices rise with the magnitude of switching costs. However, if switching costs are relatively low, as would seem to be the case with the proliferation of retail alternatives, retailers instead have an incentive to price discriminate, charging lower prices to switching customers, and higher prices to those shown to be loyal to the store (Viard 2007; Cabral 2009; Doganoglu 2010;

Pearcy 2015). Even absent the ability to price-discriminate, for instance in a setting in which common prices are set for all customers, retailers still have an incentive to invest in building market size, rather than harvesting locked-in customers (Rhodes 2014a). In fact, in a more general model that does not assume retailers are even forward-looking at all, Arie and Grieco (2014) show that sellers have an incentive to compensate customers willing to switch stores by offering lower prices, and that this compensation effect dominates the harvest effect in the short run, and the invest-effect in the long run when retailers are allowed to be forward looking. Although there are a number of potential mechanisms, the conclusion from the recent theoretical literature is the same – that switching costs need not be anti-competitive, and can drive lower equilibrium prices.

The compensation argument, however, would lead to only very small differences in equilibrium prices if consumers only need to receive their switching costs in return for moving stores. Rather, we argue that there is likely to be a substantial real option value embedded in the decision to switch. Most empirical analyses assume switching costs are either physical costs of leaving one brand and moving to another (Dube, Hitsch, and Rossi 2009), or at least "cognitive costs" associated with changing purchase patterns or investing in add-ons (Jones et al. 2000; Zauberman 2003; Murray and Haubl 2007; Osborne 2011; Liu, Derdenger and Sun 2017). In either case, the implicit assumption is that when the benefits of switching exceed the perceived costs, then switching will occur. However, we argue that this overstates the willingness to switch because switching costs are likely to also include a substantial real option value (Dixit 1992; Dixit and Pindyck 1994). That is, when there are real switching costs, and uncertainty regarding the benefits of switching to what is thought to be a lower-cost alternative, then a rational consumer will also include a real option value before making the switch. A real option is akin to a financial option in that the consumer has the right, but not the obligation, to stay with the original decision instead of switching. When the benefits of leaving are uncertain, there is a chance that the decision to stay may, in fact, be of greater value. The expected value of this option to stay represents a real option, that should have a real impact on the decision to switch. In the digital camera setting studied by Liu, Derdenger and Sun (2017) real option value arises implicitly because consumers are uncertain about the future compatibility of add-ons and as a result, switching costs are lower when these expectations are accounted for. However, Liu, Derdenger and Sun (2017) do not study equilibrium pricing. In this paper, on the other hand, we show that equilibrium pricing by retailers, who are implicitly aware of this option value, is different from the case when consumers' option value does not enter the decision.

Dube, Hitsch, and Rossi (2009) show that the magnitude of switching costs in some frequently-purchased grocery categories is sufficient to lead to lower retail prices. Indeed, the argument that loyalty generates switching costs large enough to observe is much less compelling when applied to brands as it is to stores. Nonetheless, Shy (2002), Hartmann and Viard (2008), Viard (2007), Liu, Derdenger and Sun (2017) find switching costs in non-food environments to be substantial – up to the cost of a new phone in the Israeli cell-phone market – while others find significant, but smaller switching costs in food-related contexts similar to ours (Keane 1997; Pavlidis and Ellickson 2015). Estimating switching costs in a supermarket setting, using traditional single-category scanner data, however, is problematic because it is simply not plausible that consumers choose stores based on a single category, nor are single-category switching costs likely to be as large as those associated with choosing where to shop. Because switching costs can only be inferred from assumed equilibrium behavior, there is a need for more conclusive estimates of switching costs.

We frame our empirical observations using a well-understood research data set made available by IRI, Inc. Specifically, our data describes 4 years of grocery-store visits by a panel of households in Eau Claire, WI. Panel data such as this is ideally suited for identifying store-switching costs because we observe inter-household variation in the degree to which they switch stores, the extent to which they appear to become habituated to one store, and the distances to the different stores in town. By analyzing purchases for a fixed basket of goods, we are able to study the tendencies of consumers to switch based on variations in a price-index of goods at each store. Differences in the price index at each switching occasion provide a clean estimate of the indifference point at which each household is induced to switch from one store to the next. Further, with this price data, we are able to identify how much of the price gap between stores is driving switching behavior, and how much is pure option value. By calculating the duration between switches, we are also able to obtain direct, temporal measures of the hysteretic effect of uncertainty on switching behavior. Assuming store managers are rational, and recognize consumers' real option values when setting store-level prices, we are able to test whether higher switching costs are associated with higher or lower basket prices, or the tendency of managers to "harvest or invest in" loyal customers.

We test for the price-effect of switching costs in a model of store-rivalry. In our empirical model, we assume oligopolistic retailers play an infinite series of two-period games in Bertrand-Nash rivalry (Roberts and Samuelson 1988; Perloff, Karp, and Golan 2007), where the state of competition evolves according to the number of consumers who switch, or remain loyal to their initially-preferred store. Optimal price-responses are consistent with a Markov-switching process, conditioned by the degree of loyalty expressed by consumers. By allowing retailers' price decisions to be parametric functions of the real option embedded in consumers' decisions to switch, we are able to directly test whether switching costs are pro- or anti-competitive in a typical retail food market.

We find evidence that switching costs may include a substantial real option value, which suggests that store loyalty may be the manifestation of a hysteretic response on the part of consumers to uncertainty regarding the relative prices between stores (Dixit 1992). Because stores have to compensate consumers for exercising their option to switch stores, switching costs appear to be pro-competitive, and not anti-competitive as orthodox theory would suggest. Our findings are consistent with others in the literature (e.g., Arie and Grieco 2014), but we provide a new explanation that is firmly grounded in economic theory.

Our findings contribute to both our theoretical and empirical understanding of the competitive effects of switching costs. While theoretical models of switching costs admit the possibility that they are pro-competitive, they are framed in settings where the real option embedded in consumer decision making is absent. Empirically, we examine the impact of switching costs in an environment that, arguably, has deeper implications for the economic performance of the retailing sector than previous models. While others consider switching costs for individual products, consumers do not shop intensively for items within categories (Dickson and Sawyer 1990), but do carefully consider their choice of stores (Bell and Lattin 1998). Further, the retailing sector is experiencing fundamental changes in how consumers shop, and the impact of platform-switching costs on economic performance may be central to which firms live or die, and how urban landscapes are therefore to be shaped by economic forces.

The remainder of the paper is organized as follows. In the next section, we develop a simple analytical model of real options, retailing pricing, and switching costs. With this model, we show that retailers have an incentive to offer lower prices in the presence of small switching costs, both in order to compete for a pool of loyal customers, and to compensate them for switching. In section 3, we describe an empirical model designed to test the hypotheses that follow from our theory of switching costs. We present and interpret our findings from this model in a fourth section, while a final section concludes and offers some implications for the competitiveness of the retail market.

2 Real Option Values and Loyalty

In the empirical switching-cost literature, switching costs, among other things, are interpreted as a natural consequence of state-dependence in demand (Keane 1997; Shin, et al. 2009; Liu, Derdenger and Sun 2017). That is, if the consumer reveals a significant degree of loyalty to a particular option, then price differential between the chosen option, and the next best alternative, must be sufficient to cover the implicit costs of switching. Switching costs are typically interpreted as some combination of actual financial costs and cognitive costs, or the time spent in learning about an alternative, understanding the products offered, and perhaps sampling items that are unique to the other retailer. The primary difference between a real-option model of switching costs, and one based in pure state-dependence, or loyalty, is the importance of price-volatility.

An option value arises when there are substantial fixed, or irreversible, costs, on-going uncertainty in the holder's ability to recoup those fixed costs, and a unique opportunity to exercise the option (Dixit and Pyndick 1994). Each of these conditions exists in retail grocery settings. First, the long-term volatility of retail prices is certainly not controversial - retail promotions, retailers passing through trade promotions, cost increases, or meeting competitive challenges - are all examples of price volatily over time. Second, the existence and magnitude of search costs in the retail grocery environment are documented by Mehta, Rajiv, and Srinivasan (2003). Finally, individual consumers clearly have sovereignty over their purchase decisions, so there is no market mechanism through which the real option value will be arbitraged away. In summary, the appropriate conditions exist for a purchase decision to embody a real option value, one that is significant relative to the shelf price of the product. In this section we develop a simple theoretical model of how real option values can impact the loyalty to a particular store. We show that the price-gap that should induce switching is larger with a real option than would otherwise be the case, so, when prices are volatile, loyalty will occur more often when switching costs give rise to real options.

Our mathematical derivation in a retail-price context follows that of Richards, Gomez, and Printezis (2015). Consider one component of a typical shopping basket, for instance coffee or cereal, where the production cost is largely driven by the price of the underlying commodity, and production labor. While labor is relatively stable, and predictable, the commodity cost can lead to wide variation in the costs of production. Witness the general price of groceries during the 2008 commodity boom, the 2009 bust, and the subsequent resurgence of commodity prices. Wholesale prices over this period were particularly volatile, and recent research on retail pass-through rates (Eichenbaum et al 2011; Gopinath et al 2011) suggest the same was true at retail.

Define the difference between indices of shopping-basket prices in two competing retail stores, r and s, as: $g_t = p_{st} - p_{rt}$ for period t. Note that the we define the price gap between the two stores such that a positive value is likely to increase loyalty to store r. If retail prices are uncertain, then they are expected to follow a stochastic process over time. For simplicity, we assume the relationship between the two price indices follows a Geometric Brownian Motion process:

$$dg = \mu g dt + \sigma g dz, \tag{1}$$

suppressing the time subscript, where μ is the mean drift rate, σ is the standard deviation of the process, and dz defines the Wiener increment with properties: $E(dz) = 0, E(dz^2) = dt$. Switching costs consist of opportunity costs of the consumer's time, incremental travel costs, or even cognitive processing costs that are incurred in finding a different store, locating items that are typically purchased, and making mistakes, each of which is sunk, or irretrievable after the process of switching has been completed. Define these costs generically as c_t as they are likely to vary over time.

From a consumer's perspective, her wealth is conceptualized as the present value of a stream of lifetime consumption decisions, plus the value of other investments. If the consumer's objective is to maximize the wealth, or the value, of her household by making optimal purchase decisions, the fundamental arbitrage condition that determines the value of the real option compares the value of a household that purchases from store r (V^r) with one that purchases at store s (V^s). Faced with a switching/loyalty decision at each period, the maximum value a household can attain occurs when it purchases from the store with the lowest total cost, assuming r is the current store, or:

$$V^0(g, p_{rt}, t) = \max[V^r(g, p_{rt}), V^s(p_{st}) - \lambda c_t], \tag{2}$$

where λ is the marginal cost of switching, assuming the cost is only incurred when a switch takes place. The solution to this problem provides upper and lower threshold values of g_t that constitute optimal switching points between responding and not responding to a particular value of the price gap between the two stores.

Solving for the g gives a closed-form expression for the threshold price gap as a function of the parameters of the model:

$$g = \psi \left(\lambda c(\rho - \mu) - \rho/\mu \right),$$

In the appendix we show that ψ is a function the drift rate and volatility of the price gap such that $\frac{\theta \psi}{\theta \sigma} > 0$ and is strictly greater than one. ρ is the time value of money. Because the expression on the right-side is unambiguously positive for any reasonable parameterization, this derivation shows that the presence of a real option value in store-switching behavior causes consumers to switch at larger price-differentials than would otherwise be the case. In other words, the price gap g, which is the difference in prices that would induce a consumer to switch from one firm to another is increasing in σ, i.e., $\frac{\theta g}{\theta \sigma} > 0$. Expressed in terms of the rate of store-switching, our real option model implies a hysteretic effect, or a substantial delay, relative to the case where a real option is either not present, or is ignored. Because real option values rise in the volatility of prices, our theory implies that state-dependence rises in the volatility of prices, however it does not imply that the choice probability rises in price-volatility, as the variability of prices is likely to be regarded as a negative attribute, but that the persistence in demand for a particular option rises in the level of price volatility. Therefore, in the empirical model below, we test the real-option theory of store-loyalty by allowing the state dependence in demand to depend on a "mere loyalty" effect that is due simply to purchase history, and a second "real option" effect that is due to price volatility.

3 Econometric Model

Our model of equilibrium pricing assumes that oligopoly retailers compete in prices over an infinite horizon. That is, our equilibrium concept is a Markov Perfect Equilibrium (MPE) in prices, conditional on the effect of consumer switching costs on market shares. Consistent with the literature on this topic, retailers are unable to price discriminate among loyal and switching consumers, so set prices that manage the harvest and invest incentives simultaneously. As a result, we capture the long-term equilibrium between the two retailers in our sample market as if they are conscious of the two market segments co-existing. In a two-retailer environment, this assumption leads to a model of equilibrium pricing that is both tractable and descriptive of the retail environment, as the majority of consumers consider one of the two stores for their weekly grocery needs. In this setting, if switching costs are pro-competitive in effect, then we expect to find equilibrium prices lower with higher switching costs, and vice versa.

We first test this hypothesis informally by looking at reduced form evidence in Section 4.2 and then we explore the question more formally using a set of counter-factual numerical experiments. Namely, we test for the competitive effects of state dependence through mere loyalty by comparing equilibrium prices with and without loyalty. If loyalty-driven switching costs are associated with lower equilibrium prices, then it is reasonable to conclude that switching costs are pro-competitive. Then, we examine the same question with the addition of switching costs that derive from the real-option effect. If we remove the real option associated with moving from one store to another, and find that prices rise as a result, then we conclude that this element of switching costs is pro-competitive as well. In either case, we are lead to a dramatically different implication from what is currently assumed to be true: Locking in customers is a way to ensure higher prices in equilibrium. We test this hypothesis using the econometric model developed in the next section.

Our empirical model of price competition is conditioned on a model of store choice, such that the state of the system is fully described by the number of consumers switching from one store to the other. The conceptual model that underlies this approach is a two-stage game between retailers, and retailers and consumers. In the first stage, retailers set prices for items in specific categories consistent with an overall competitive basket-price target. The assumption that retailers compete in basket-level prices reflects the notion that retailers cannot compete category-by-category, but rather choose prices in order to meet a store-level goal for price-competitiveness. In the second stage, consumers see the prices posted by each retailer, and make a store choice with these prices, and other attributes of the stores that we include in our model, in mind. To implement this model, we first estimate a nested logit model in which consumers choose to purchase a basket of goods from one of two stores, where prices form part of the attractiveness of each store. We use this model to estimate the number of consumers who switch stores each period, and their cost of switching, conditional on their revealed loyalty to each store. We then estimate a model of price competition between the stores in which the optimal vector of prices chosen by each store is a function of their customers' cost of switching. We then simulate this equilibrium in order to demonstrate the potential profitability of pricing with switching behavior in mind.

3.1 Empirical Model of Category and Store Choice

Our empirical model of store-switching costs assumes a hierarchical structure of category and store choice. That is, we follow Briesch, Chintagunta and Fox (2009) in assuming that consumers follow a three-step process in choosing a store: They first decide which categories they need to re-stock, they then form an estimate of the utility attainable from each store, based on the attributes of the store and the items it offers, and then choose the store in order to maximize utility.[1] Modeling category-choice as part of a store-choice model is important, because doing so allows us to control for how category-level attractiveness affects the choice of shopping destination.

We begin our model development at the category-choice stage, and then describe a model of store choice. Utility from category-and-store choice is additive, so express the joint utility of category-and-store purchase as:

$$U_{jhrt} = V_{jht|r} + W_{hrt} + \varepsilon_{jhrt}, \tag{3}$$

[1] Note that this assumption does not carry any behavioral implications. A nested logit model merely specifies a mathematical correlation among the category and store choices and does not imply that they are sequential in any way. We are agnostic on that point.

for the choice of category j in retailer r by household h on purchase occasion t, where ε_{jht} is a random error term that captures heterogeneity in preferences across households for both categories and retailers (which we decompose into category- and store-based components below) and is assumed to be Generalized Extreme Value (GEV) distributed, $V_{jht|r}$ is the deterministic component of utility associated with the choice of category j, conditional on the choice of retailer r, and W_{htr} is the utility derived from purchasing at retailer r.[2] Based on the distributional assumptions for ε_{jht},we then derive the joint probability of choosing categories and stores such that: $\Pr(c_j, r) = \Pr(c_j|r)\Pr(r)$, where $\Pr(r)$ is the probability of choosing retailer r, $\Pr(c_j, r)$ is the joint probability of choosing category j and retailer r and $\Pr(c_j|r)$ is the conditional probability of choosing category j, given that r has been chosen.

We introduce state-dependence in our demand model through the store-choice stage, driven by habituation and switching costs, as our focus is on store-loyalty and the cost of switching between stores.[3] In this way, we model store choice in a tractable, nested-logit framework that follows others in this literature (Bucklin and Lattin 1992; Bell and Lattin 1998; Briesch, Chintagunta and Fox 2009) and yet contains dynamic elements that are likely to be essential to competition among stores that interact in the same market on a day-to-day basis.

3.2 Category Choice

The conditional choice of category is driven by both preference and need. While some households may prefer carbonated soft-drinks to fruit juice, once this preference is established, the probability of category purchase (purchase incidence) on any given visit to the store depends upon the household's rate of consumption, the quantity they purchased at the last visit, and the amount of inventory on hand (Bell, Ho, and Tang 1998; Bell and Lattin 1998;

[2] Train (2003) describes the difference between the logit and nested-logit in terms of the properties of each ε_{jht} draw. For the simple logit, each ε_{jht} is an independent draw from a univariate extreme-value distribution, while each marginal distribution for the ε_{jht} in a nested logit is an independent draw from a univariate extreme-value distribution.

[3] Osborne (2011) develops an empirical model in which state-dependence at the item-level depends on both switching costs and learning, but neither of our focal stores is new to this market, so the notion of learning is not plausible in our setting.

Briesch, Chintagunta, and Fox 2009). Category incidence also depends on heterogeneous household preferences, both associated with observed attributes of the household, and unobserved elements that vary at the household level. We separate these variables into sets of need-based ($\mathbf{X}_{jht}$) and household-based ($\mathbf{Z}_h$) specific factors such that the utility obtained from purchasing in category j, conditional on the choice of retailer r is written as:

$$V_{jht|r} = \alpha_{jh}^{r} + \beta_h \mathbf{X}_{jht} + \gamma \mathbf{Z}_h, \tag{4}$$

where α_{jh}^{r} is a fixed category-effect variable that accounts for variation in category preferences among households (which we allow to vary by retailer), while the vector of household factors ($\mathbf{Z}_h$) consists of the size of the household (HH_h), the age of the household head (AGE_h), income (INC_h), and educational attainment (ED_h). Our vector of need-based variables at the category level ($\mathbf{X}_{jrt}$) is intended to capture the amount on hand at any purchase occasion, and the probability that current stockpiles are likely to be depleted soon. To this end, we calculate the average consumption rate over a 6-month initialization period at the household level (CR_{jht}), the amount purchased on the previous shopping trip (LQ_{jht}), and the number of days since items in the category were last purchased (ITT_{jht}). With this specification, we control for variation in preferences, and revealed consumption behavior that are likely to explain why a household chooses to purchase from a particular category, regardless of where they shop (Briesch, Chintagunta, and Fox 2009). Although the IRI data contains a large number of potential explanatory variables, this set is both well-accepted in the literature, and is likely to capture most of the observed heterogeneity in category choice.

Identification of state dependence requires that other sources of unobserved heterogeneity are appropriately controlled for (Keane 1997; Train 2003; Osborne 2011). Indeed, state dependence and choice heterogeneity are observationally equivalent if not measured precisely, and will pass through to the store-choice level if not addressed in the conditional, category-choice stage. We account for unobserved heterogeneity at the category level by allowing for random variation of the category preference (α_{jh}^{r}) and need-based (β_h) parameters, defining each to be normally distributed so that:

$$
\begin{aligned}
\alpha^r_{jh} &= \alpha^r_{0j} + \alpha^r_{1j}\mu_{1j} \\
\beta_{hi} &= \beta_{0i} + \beta_{1i}\mu_2,
\end{aligned} \tag{5}
$$

where μ_{1j} and μ_2 are each standard normal variates and α^r_{0j} and β_0 are the mean category preference terms for category j in retailer r and for need-based element i, respectively. Consistent with the GEV assumption described above, the conditional choice of each category j by household h is a binary logit written as:

$$
P(c_{jh}|r) = \frac{\exp(\alpha^r_{jh} + \beta_h \mathbf{X}_{jht} + \gamma \mathbf{Z}_h)}{1 + \exp(\alpha^r_{jh} + \beta_h \mathbf{X}_{jht} + \gamma \mathbf{Z}_h)}, \tag{6}
$$

where consumers are assumed to make $j = 1, 2, ..., J$ decisions on each purchase occasion. The total value of all J decisions, however, affects the choice of store as described in the second, store-choice stage of the model.

3.3 Store Choice

Store choice depends upon attributes of the store, and both observable and unobservable attributes of the consumer that may mean that one store is more desirable than another. Bell, Ho, and Tang (1998) differentiate between fixed costs of visiting a particular store – those that do not vary with the amount purchased – and the variable costs, or those that depend on how much the consumer buys. Variable costs include measures of the relative cost of shopping at each store, typically captured by a price index (Bell, Ho, and Tang 1998), elements of the marketing-mix, the assortment available at each store (Briesch, Chintagunta, and Fox 2009), and the relative attractiveness of purchasing a basket of items from each store (Bell and Lattin 1998). Fixed costs include the distance between the household and each store, as well as measures of state dependence that capture opportunity costs of departing from habits or familiar patterns of behavior (Bell and Lattin 1998; Briesch, Chintagunta, and Fox 2009). In fact, we depart from others in this literature by including a structural model of state dependence, driven by the real option embedded in consumers' store-choice decisions (Ching, Erdem, and Keane 2014).

* to here - re-derive store-choice model using CEK method.

Consumers' store-choice decisions are inherently dynamic because the choice today affects the information set that determines the choice made in the next, and all future, periods. We capture all elements of store choice in writing the indirect utility from choosing store r as:

$$W_{ht}^{r}(r_{ht}) = \delta_h \mathbf{Y}_{hr} + \phi_h I^h(r=1) + \lambda_h BA_{hr}, \tag{7}$$

for a set of store variables $\mathbf{Y}_{hr}$ that includes the distance, measured in Euclidean terms, between household h and store r (DT_{hr}), a measure of the depth of assortment across all categories in store (VR_r), calculated as the total number of universal product codes (UPCs) offered each week in store r, a sales-weighted price index defined across all categories in store r (PR_r), an index of promotional activity (PM_r), defined as a sales-weighted average of the percentage of items purchased on promotion during a given week, a metric of featuring defined in the same way (FT_r), a control-function designed to address the endogeneity of pricing decisions (CT_r), a household-level indicator of store-loyalty, our measure of state-dependence (I^h), defined as 1 when the household is loyal to a particular store, or when the household visits the same store on a purchase occasion as they did on the previous trip, and the "basket attractiveness" (Bell and Lattin 1998) term (BA_{hr}) for store r and household h. The concept of basket attractiveness captures the relative attractiveness of the store as measured by the expected utility from a basket of groceries estimated in the conditional basket choice model, written as: $BA_{hr} = \sum_{j \in J_r} [\log(1 + \exp((\alpha_{jh}^{r} + \beta_h \mathbf{X}_{jht} + \gamma \mathbf{Z}_h)/\lambda_h))]$.[4] As in the category-choice model, we account for unobserved heterogeneity in store preference ($\delta_0 hr$), store-loyalty (ψ_h), value of the entire basket (λ_h), and the other store-specific attribute variables by allowing each to be randomly distributed over households, using the same distributional assumption

[4] Basket attractiveness is analogous to the inclusive value term of more usual nested logit models. Defining the inclusive value term as a sum of the utility obtained from each category assumes that the utilities are independent. As Bell and Lattin (1998) explain, this assumption is valid for all but categories that are obvious complements, such as pasta and sauce or hot dogs and ketchup.

as in the category-choice model.

In this model of store-choice utility, however, the indicator of store-loyalty is likely to be endogenous as unobservables in the utility expression are almost certainly correlated with the binary measure of loyalty. In fact, the theoretical model of switching costs described above (and more formally in the appendix) implies that the importance of loyalty should depend on the price-differential between the two stores, as well as the volatility of the difference in prices. We use our model of the real option embedded in store switching to suggest an instrumental-variables approach to accounting for the endogeneity of loyalty (Davidson and MacKinnon 2004) as a function of the price-differential between store r and store s on each purchase occasion, as well as the standard deviation of the difference in prices and promotional activity in each store. We then retain the Inverse Mills' Ratio (IMR), and include its value as an argument in the store-choice model as an instrument for the endogenous loyalty measure. With this correction, the nested logit will produce unbiased estimates of the effect of loyalty, and hence switching costs, on store choice. Moreover, using the counterfactual simulations described in the pricing model below, this approach also allows us to test for the statistical significance of real option values on the cost of store-switching.

To complete the nested logit model structure, the GEV assumption for ε_{jht} implies that the choice among stores takes a logit form, written as:

$$\Pr(r) = \frac{\exp(\delta_h \mathbf{Y}_{hr} + \psi_h IMR^h + \lambda_h BA_{hr})}{\sum_{s \in R} \exp(\delta_h \mathbf{Y}_{hs} + \psi_h IMR^h + \lambda_h BA_{hs})}, \tag{8}$$

with the basket-attractiveness term as defined above. With estimates of the store choice model, we can then calculate the joint probability of each category-and-store choice by integrating over the distributions of heterogeneity as in:

$$\Pr(c_j, r) = \Pr(c_j|r)\Pr(r) = \int\int\int\int\int \left(\frac{\exp(\alpha_j^r + \beta_h \mathbf{X}_{jr} + \gamma \mathbf{Z}_h)}{\sum_{s \in R} [\exp(\alpha_j^r + \beta_h \mathbf{X}_s + \gamma \mathbf{Z}_h)]} \right)$$
$$\left(\frac{\exp(\delta_h \mathbf{Y}_{hr} + \psi_h IMR^h + \lambda_h BA_{hr})}{\sum_{s \in R} \exp(\delta_h \mathbf{Y}_{hr} + \psi_h IMR^h + \lambda_h BA_{hr})} \right) d\mu_1 d\mu_{2j} d\mu_\alpha d\mu_\phi d\mu_\lambda. \tag{9}$$

where the $d\mu_k$ terms refer to each of the random coefficients in the category, and store-choice specifications. Although the nested-logit model has a closed form in the absence of unobserved heterogeneity, the fact that the preference-distributions for each household are not observed means that we have to integrate over the distribution of heterogeneity for each household. Because of this fact, we estimate using simulated maximum likelihood as described below.

In the nested store-choice model, the key parameters that determine the importance of inter-store choice are the store-loyalty estimates. In the store-choice model, the importance of loyalty is given by the ψ_h estimates in equation (7). If this parameter is significantly greater than zero, then state dependence is an important determinant of store choice. More importantly, the significance of this parameter suggests that retail-stores may compete in a dynamic game for market share, inducing consumers to overcome their inherent loyalties by either reducing basket prices, or harvesting their own loyal shoppers by maintaining high basket prices.

Testing whether this is the case, however, requires that we use our model of individual-store-choice to describe market dynamics. That is, the model of store-choice is estimated at the individual level, but our structural approach to pricing described in the next section requires market-level information. Therefore, we aggregate the choice-probability model in (9) over the distribution of sample households on a weekly basis in order to arrive at a measure of weekly market share for each firm, j (S_{jt}). Because market share is a function of the proportion of consumers loyal to each store, the extent of loyalty, and switching, provides the dynamics required to capture any intertemporal pricing behavior designed to induce switching, or to harvest each store's own loyal consumers.

In the demand model, store-level prices are likely to be endogenous. Therefore, we estimate the demand model using a control-function approach (Petrin and Train 2009), using as instruments key input prices for each of our included categories, and input price indices designed to capture variation in retailing costs. Input prices at a category-level are likely to be valid instruments because acquisition costs are correlated with output prices according to fundamental principles of economics, and yet mean independent of any unobservable demand factors. While we cannot assess the empirical validity of this logic in a structural way, a first-stage regression of our instruments on the aggregate price indices can suggest whether they are weak in the sense of Staiger and Stock (1997). For the two stores in our sample data, the first-stage IV regression produces an F-statistic of $11,444.1$, which is much larger than the threshold value of 10.0 suggested by Staiger and Stock (1997), so our instruments are not weak.[5]

3.4 Model of Strategic Pricing

Our model of strategic pricing follows the logic of Roberts and Samuelson (1988) and Perloff, Karp, and Golan (PKG, 2007). In this model, two rival firms compete in price, conditioned on the state of the market, where the state evolves according to the number of consumers who are willing to switch stores each period. Switching, in turn, is governed by a Markov switching process described by the loyalty dynamic above. Our model is inherently simpler than that described in PKG as firms compete only in price, but is complicated by the fact that switching does not follow the simple linear accumulation rule described in their model.

Optimal intertemporal price decisions are derived under the assumption that retailers play a general Nash game in prices. Although retail competition is typically a multi-faceted game, we focus on competition in basket-prices in order to capture the strategic nature of pricing, and its relationship with store-switching costs. At a store-level, most other potentially-strategic variables are arguably pre-determined by the size of the store footprint, or the

[5] Detailed estimates from the IV regression are available from the authors.

decision on where to locate the store. Therefore, price is the one variable that can be varied over a typical retail decision-cycle (one week) to attract customers away from a competing store. Implicitly, therefore, we assume that store-level pricing is the primary means by which managers are able to overcome loyalty to another store, and induce switching behavior in intertemporal competition. We model dynamic competition among stores following Roberts and Samuelson (1988) and Perloff, Karp, and Golan (2007) by specifying a simple two-period dynamic optimization model in order to capture the intertemporal dimension of interfirm reactions in price. As Perloff, Karp, and Golan (2007) explain, the model is a "hybrid" approach that does not oversimplify the nature of competition between the firms as an open-loop solution would imply, but is empirically tractable, unlike a Markov-perfect model. By using this approach, we capture strategic reactions over time in a way that neither assumes away the interactions we hope to test, nor complicates estimation in a way that excludes all but the simplest of structural frameworks. Because strategic interaction as a source of dynamic supply behavior remains a testable hypothesis, we retain the possibility that reactions can also occur in the current period (are *ex ante* rather than *ex post*) in order to test for the appropriate form of the game. In this way, we present a more general treatment of interfirm price rivalry.

Given the assumptions regarding the market-demand for each store outlined above, the two-period profit function faced by firm j is:[6]

$$V_{j0} = (\pi_{jt} + \rho\pi_{j,t+1}) = (p_{jt} - c_{jt})S_{jt}M + \rho\left((p_{j,t+1} - c_{j,t+1})S_{j,t+1}M\right), \tag{10}$$

where ρ is a discount factor, M is the size of the entire market, and c_j is the marginal cost of production for firm j. The marginal cost of production is assumed to be derived from a Generalized Leontief (GL) unit-cost function, C_j. A GL cost specification is chosen because it is flexible (meaning that it is an approximation to an arbitrary functional form), it is inherently homogeneous in prices without normalization, it is affine in output without

[6]We develop the model in general notation for $i = 1, 2, ..., I$ firms, but note that the model simplifies considerably in our two-retailer case.

further restriction, and it imposes convexity in output, while concavity in prices, symmetry, and monotonicity can be maintained and tested. Total production and marketing costs are a function of the primary inputs to grocery production at the category level, and in retail marketing so the cost of producing one unit of output is:

$$c_j(\mathbf{v}) = \sum_{k=1}^{K} \tau_{jk} v_k + \sum_{k=1}^{K}\sum_{l=1}^{L} \tau_{jkl} (v_k v_l)^{1/2} + \mu_j, \tag{11}$$

where v_k is a vector of input prices and μ_j is an iid random error term. With this unit-cost function, the marginal cost of grocery marketing is also Generalized Leontief, so retains the attributes described above.

Our model of dynamic price competition admits a range of possible behaviors. Roberts and Samuelson (1988) describe a firm that does not anticipate the reactions of its rivals as "naive" and one that does as "sophisticated." We extend this concept to differentiate between *ex ante* expectations of current-period reactions and *ex post* responses by rivals. In this way, we are able to test for which, or both, represents an appropriate description of the game played by rival supermarkets. Assuming a general Nash equilibrium, the first order conditions to the problem defined in (10) provide optimal response equations in basket-level prices. The necessary condition for an optimal price response is:

$$\begin{aligned}\frac{\partial V_{j0}}{\partial p_{jt}} &= S_{jt} M + (p_{jt} - c_{jt}) M \left(\frac{\partial S_{jt}}{\partial p_{jt}} + \sum_{i=1}^{I} \frac{\partial S_{jt}}{\partial p_{it}} \frac{\partial p_{it}}{\partial p_{jt}} \right) \\ &+ \rho (p_{j,t+1} - c_{j,t+1}) M \left(\frac{\partial S_{j,t+1}}{\partial p_{jt}} + \sum_{i=1}^{I} \frac{\partial S_{j,t+1}}{\partial p_{i,t+1}} \frac{\partial p_{i,t+1}}{\partial p_{jt}} \right) = 0,\end{aligned} \tag{12}$$

where i indexes rival firms, and we define $\phi_{ij} = \partial p_{it}/p_{jt}$ as the expected current-period response of firm i to a change in the price of firm j, and $\omega_{ij} = \partial p_{i,t+1}/p_{jt}$ as the expected intertemporal response. Substituting the expressions for nested-logit share derivatives: $\partial S_j/\partial p_j = \frac{\alpha}{\lambda} S_j(1 - (1-\lambda)S_{j|r} - \lambda S_j)$ and $\partial S_j/\partial p_i = -\frac{\alpha}{\lambda} S_j((1-\lambda)S_{i|-r} + \lambda S_i)$ into (12), aggregating the household-level estimates to the market-level, and solving for the retail margin in period t gives an estimable equation for the price-response of firm i in terms of the utility function parameters, the expected margin and expected competitor reactions:

$$(p_{j,t} - c_{j,t}) = \left(-\frac{\alpha}{\lambda}(1 - (1-\lambda)S_{j|r,t} - \lambda S_{j,t}) + \sum_{i=1}^{I} \frac{\alpha}{\lambda}((1-\lambda)S_{i|-r,t} + \lambda S_{i,t})\phi_{ij} \right)^{-1} \tag{13}$$
$$- \rho \left(\frac{(p_{j,t+1} - c_{j,t+1})(\sum_{i=1}^{I} S_{i.t+1} S_{j,t+1}\omega_{ij})}{S_{j,t}(1 - (1-\lambda)S_{j|r,t} - \lambda S_{j,t}) - \sum_{i=1}^{I} S_{j,t}((1-\lambda)S_{i|r,t} + \lambda S_{i,t})\phi_{ij}} \right) \forall j,$$

where $\phi_{ij} \neq 0$ indicates a departure from the Bertrand-Nash pricing rule that is typically assumed in existing retail pricing models at the store level. We then estimate the marginal cost and strategic pricing equation in (13), after substituting the parameters from the demand model, and observed values for all market shares and prices. In this way, we characterize each firm's optimal price response, conditional on both the structure of demand, and the optimal response of its rival. While the parameters of the supply model are of some inherent interest, we cannot yet test directly for whether the extent of loyalty revealed by the demand model leads to higher, or lower, prices in equilibrium.

For that purpose, we conduct a number of counterfactual simulations of the supply model in (13) under different assumptions regarding the importance of switching-costs, and examine whether equilibrium prices at the store level are higher, or lower, than the benchmark scenario of zero switching costs. Specifically, we compare equilibrium prices under three scenarios: (1) the base-case solution with the estimated degree of state-dependence, and switching costs, (2) no switching costs due to mere loyalty, or state dependence that is due only to consumers' tendency to shop at the same store as they did on the last shopping trip, and (3) no switching costs due to the absence of real options mediating the state-dependent effect that is due to mere loyalty. With this approach, we not only answer the central question as to whether switching costs are pro- or anti-competitive, but are able to estimate the relative effect of stores competing for loyal customers versus compensating consumers for exercising the real option associated with switching stores.[7]

[7] By estimating the demand for each store with data from the 10 most popular categories, we avoid the problem of secondary-shopping trips that would otherwise confound our definition of loyalty. In our model, secondary shopping trips are only used for purchasing minor categories, so consumers can remain loyal to a primary store for most of their weekly shopping needs, and still face substantial switching costs if they purchase their primary needs at another store.

4 Data Summary and Empirical Observations

4.1 Descriptive Statistics

Our data are from the IRI Academic Data Set (Bronnenberg, Kruger, and Mela 2008) in the Eau Claire, WI market over a period of 4 years (2008 - 2011). In order to ensure that we are able to identify store loyalty, and switching behavior, we focus on the households with at least 800 category-store choice observations, resulting in data from 2,737 households over the entire sample period. For each household in our sample, we record their purchases from the top 10 most-frequently purchased categories, from each of 2 focal stores in the Eau Claire market, as well as purchases from these categories made in any other outlet. These non-focal-store purchases serve as our outside option. For each household and purchase occasion, we have data on the specific item purchased, the amount that was purchased, the price paid, and any marketing activity (price-promotion, featuring, display) that the household may have taken advantage of or been exposed to. The IRI household panel data set also provides a rich set of demographic variables, including the age of the household head, household income, the age of each household member, and the number of children present in the household. Data from items that were not purchased on each shopping trip were imputed by using store-level data that is also included as part of the IRI Academic Data Set.

The set of items households purchase from shopping trip to shopping trip is likely to vary. If store choice is driven by advertised specials on specific items, resulting in one-time trips to an alternative store, this is not switching, but rather "cherry picking" the best deals. While cherry picking is important and prevalent (Fox and Hoch 2005), it does not describe the type of behavior that we focus on. Indeed, our definition of habituation and store loyalty still admits some cherry-picking behavior as a realistic element of shopping behavior. Our definition of loyalty follows Dube, Hitsch, and Rossi (2009) in that a household is defined as loyal if it visits the same store on the current purchase occasion as they did on the previous shopping trip. They may switch back subsequently, but we assume that back-switching involves less cost because the store layout, quality, and other attributes will be well-understood. According to this definition, therefore, we need to be explicit about what constitutes a "basket" of purchases for a household.

In order to capture the relative attractiveness of the products offered in each store, we model households' category incidence in a manner similar to Bell and Lattin (1998); Bell, Ho, and Tang (1998), and Briesch, Chintagunta and Fox (2009). In this way, we implicitly assume that each household contemplates visiting one of the focal stores (or the outside option) in order to fulfill a primary shopping need from one of 10 most-frequently purchased grocery categories. We choose these categories based on purchase frequency and household penetration specific to the population of households in the IRI panel (these data are reported in Bronnenberg, Kruger, and Mela 2008). Specifically, the top 10 categories include: Milk, cereal, carbonated soda, salty snacks, frozen entrees, soup, toilet paper, yogurt, beer, and coffee. Together, these categories account for over 65% of weekly purchases from our sample households. While our categories may be deemed sub-categories according to more usual definitions (i.e., milk is a sub-category of dairy, and soda is a sub-category of beverages) our definition is necessary in order to track prices at the appropriate level of aggregation, without becoming too specific and tracking purchases of brands, and mistaking brand loyalty for store loyalty. For each category, we include private labels among the products purchased in order to capture as much store-specific category attractiveness as possible. Table 1 provides a summary of the category and store-level data.

[table 1 in here]

Identifying switching behavior, and switching costs, requires, of course, evidence of switching behavior in our data, and that behavior must by driven an exogenous source of variation. This latter point is important because if switching behavior were purely endogenous, or driven by normal variation in prices, then there would be no way to identify the cost of switching as opposed to simply responding to price variation. In this regard, we follow Dube, Hitsch, and Rossi (2009) in assuming switching behavior is driven by promotional activity, which is largely driven by retailers' responses to trade promotions, and, as such, is exogenous to retailers' daily pricing decisions (Gomez, Rao, and McLaughlin 2013). Elements of store-attractiveness that can be varied in this way are termed "variable" factors by Bell, Ho, and Tang (1998) as they can be changed over a relatively short timeframe by store management. Our variable factors are defined as category-weighted averages of any price promotion, display, or feature activity that the household may have been exposed to on each trip to the store. Assortment depth is another variable factor, which we capture by calculating the total number of SKUs on offer across at all 10 categories, on a weekly basis. We expect that a deeper assortment will attract a greater number of households, regardless of any price-induced switching behavior.

Switching behavior may also be due to store-level variables that are independent of any temporary marketing activity. Therefore, we define several variables that capture the relative attractiveness of each store, categorized as "fixed" elements of a store, or of consumer behavior. Following Bell, Ho, and Tang (1998), our primary fixed factors include a measure of distance between each household-and-store pair, and a measure of store loyalty that captures the primary dynamic in our model of store choice. Distance is calculated from the center of each household's ZIP code to the center of the store ZIP code. Although we recognize that ZIP codes are relatively large geographic aggregates, in the absence of exact street addresses they are likely to be at least broadly reflective of the distances to each store. Loyalty is calculated in a way that captures the intertemporal nature of consumers' tendency to switch between stores. Unlike Briesch, Chintagunta, and Fox (2009), who define store loyalty as a static metric based on visit-shares in an initialization period, we measure loyalty in terms of the last-purchase made as in Dube, Hitsch, and Rossi (2009). In order to capture the evolution of loyalty as a dynamic measure, we define a household as loyal to a particular store if their last purchase was made in that store, and the household remains loyal to that store until a switching event occurs. While this is a fragile notion of loyalty relative to more static measures, it captures the fact that our households evidently regard the stores as relatively close substitutes (table 1) so loyalty is, in aggregate, not absolute.

Prior to presenting the results of the structural model above, we examine the data for reduced-form evidence of our maintained hypothesis, namely that store loyalty is a hysteretic effect, largely driven by price volatility.

4.2 Reduced Form Evidence

The theoretical model described in Section 2 (and more formally laid out in the Appendix) implies that the importance of loyalty should depend on the price-differential between the two stores, as well as the volatility of the difference in prices. Thus we start our empirical exploration by informally exploring whether there is any reduced form evidence for the relationship between these variables. We then estimate a more formal structural model and discuss the full structural results in the next subsecion.

To achieve this, we estimate a simple probit model in which loyalty is written as a function of the ratio of the price index of the competing to the own-store (defined such that higher values are likely to be associated with a higher probability of a consumer being loyal to the own-store)[8], a measure of price volatility calculated as a 4-week moving standard deviation of the own-price, a set of demographic variables (age, education, income, and family size), and marketing-mix activity in the own-store (assumed to be exogenous as it is largely driven by trade promotions)

[table 2 in here]

Table 2 reports the results. As expected, lower prices and more promotional activity (except in the case of featuring in store 2) are associated with a higher degree of loyalty. More important for current purposes, however, the estimates in this table show that price volatility is positively related to loyalty, regardless of the store, although the point estimate is only significant at a 6% level (with a one-tailed test) in Store 2. While there may be other reasons for this finding, this preliminary reduced form evidence is consistent with the notion advanced in our theoretical model, namely that a larger real option value generates a hysteretic effect in store choice, which is manifest in greater store loyalty. Each dollar increase

[8] We also estimated the specifications with absolute price differences and found almost identical results.

in the standard deviation of prices is associated with an approximately 12% increase in the probability a household is loyal to store 1, and a 5% increase in the probability of store 2 loyalty. Finally, we find that older, larger, higher income, and better-educated households tend to be more loyal to Store 1, but less loyal to Store 2. Clearly, each retail chain differs in their ability to drive loyalty to particular market segments.

Overall, we find convincing preliminary evidence consistent with our theory that loyalty rises in the volatility of prices, suggesting that store loyalty is driven not only by state-dependence but also by hysteretic effect arising from price volatility. Next, we present the results of our structural model of utility maximization and pricing.

4.3 Structural Results and Discussion

In this section, we present and discuss the estimates from our structural model of store-choice and strategic pricing. We first present the results from the model of demand, and establish a preferred specification for how state-dependence enters into the model. We next present our estimates from the pricing, or supply-side, of the model, and a set of counterfactual simulations with the pricing model that are intended to test our core hypotheses regarding the effect of switching costs on equilibrium prices, and the role of real options as an alternative explanation for how switching costs arise. We conclude with a discussion of the implications for both the competitiveness of food retailing, and for switching behavior more generally.

Both levels of the nested-logit model of store choice are estimated together, but we present the estimates in 2 separate tables. The estimates in table 3 report our findings from the category-choice stage, and we show estimates from a fixed-coefficient version of the model as well as a version that allows for random parameters in the key household need variables (lagged volume, consumption rate, and interpurchase time). We first use a likelihood-ratio (LR) test to evaluate which specification, fixed- or random-coefficient, provides a better fit to the data. With 12 random coefficients, the critical Chi-square value is 21.026, while the estimated LR statistic is 103.48, so we reject the fixed-coeffiicient version in favor of the random coefficient model. This outcome was to be expected because household need is likely to include a substantial amount of unobserved heterogeneity, even after controlling for the set of observed household attributes shown in the results table. As expected, consumption rate and interpurchase time have strong, statistically significant effects on the probability a household purchases one of the 10 categories on each visit to the store. Controlling for consumption rate, and inter-purchase time, lagged purchase-quantities can have one of two effects on purchase incidence. First, a stockpiling effect will have a negative effect on purchase incidence, because if a household purchased a large amount on the last visit, relatively little is needed on the current trip (Briesch, Chintagunta, and Fox 2009). Second, lagged volume may also capture a habitual-purchase effect as households that purchase large quantities on one visit are more likely to choose the same category on the next trip to the store because they are heavy, habitual purchasers. Our estimates reflect this latter effect as the coefficient on lagged volume is positive, and strongly significant. Heavy users tend to purchase more on each visit, and purchase more frequently. The fact that this variable is picking up some of the dynamics in category demand, moreover, means that our estimate of state-dependence at the store level, will be more conservatively estimated.

[table 3 in here]

Our focus in this paper is on store-choice, so the estimates in model 4 are of more immediate interest than the purchase-incidence results reported in table 3. The parameter that links the purchase-incidence model to store choice, λ, suggests that consumers place a relatively small, yet highly statistically significant, value on the attractiveness of the entire basket offered by each store. Given that our choice of categories spans those typically available at each store, and the categories vary only slightly in content between stores, this outcome is perhaps to be expected. Among the other variables that affect store choice, we find that the price index is highly significant, and negatively-signed, as expected. Clearly, consumers form expectations of what an average "basket price" would be at either store in making their decision as to where to shop. Moreover, the importance of price in choosing between stores suggests that the price-gap between the two stores should be an important driver of switching behavior, and that a real option is likely to arise. We also find that variety is an important component of store choice, consistent with Hoch, Bradlow, and Wansink (1999), Briesch, Chintagunta, and Fox (2009), and Richards and Hamilton (2016). When consumers are heterogeneous in their preference for food and beverages, a deeper assortment is likely to appeal to a greater number of consumers. Further, driving distance is a key component of fixed shopping cost (Bell, Ho, and Tang 1998), so greater distance from the household to the store reduces the probability that a consumer chooses a particular destination. Among the other marketing-mix variables, we find that the index of promotion has a negative effect when featuring activity is also included in the model. This result suggests that featuring is a more important driver of traffic than price-promotion, and largely achieve the same, overlapping, objective.

[table 4 in here]

The most important element of the store choice model, however, both conceptually and for current purposes, is loyalty. Finding that loyalty is, in a statistical sense, the most important driver of store choice is consistent with most of the literature on store choice that we have cited, and with retail practice (Reichheld 2003). Store managers understand that consumers tend to shop habitually, so that once a customer is attracted to the store, they are reluctant to leave. Despite the statistical importance of loyalty, our summary statistics above show that switching is a feature of our data, so there must be a price-gap between the stores that exceeds both the implicit and explicit costs of switching. In fact, the estimates reported in table 4 imply that total switching cost, at the shopping-basket level, is approximately $21.37, or 14.22% of the total cost of a basket of groceries.[9] Based on this estimate, switching costs are not only statistically significant, but substantial in an economic sense.

Recall that, in our model, loyalty is endogenous in the sense that it is determined simultaneous to the consumers' category- and store-choice decisions. Thus we use instrumental

[9] We calculate this estimate by finding the "willingness-to-pay" for loyalty, or the loyalty parameter divided by the marginal utility of income estimate.

variable approach[10]. We then estimate the inter-temporal pricing model in (13) using the store-choice and loyalty estimates to calculate each element of the model. Because the pricing model is highly non-linear, we estimate this final element of our structural model using non-linear least squares, both with fixed parameters, and a random parameter version to control for unobserved heterogeneity in the key conduct parameters, ϕ and ω. In this model, we estimate the marginal cost of producing and delivering our hypothetical shopping-basket of items, and the short- (ϕ) and long- (ω) term competitive conduct of the two stores. Our minimal set of input prices consists of an index of retail wages, and of business services.[11] In terms of retailing cost, the estimates in table 5 show that they are dominated by business services, which include advertising, insurance, and management. Because these price indices tend to be highly correlated with other input prices, including business services makes the others statistically insignificant. After controlling for weekly variation in selling costs, we find that the short-run competitive conduct is only slightly more competitive than the maintained Bertrand-Nash assumption ($\phi = 1$) as $\phi = 0.95$ from the best-fitting (fixed-coefficient) model. Although the departure from Bertrand-Nash conduct is statistically significant ($t = 59.94$), this estimate suggests that the Bertrand-Nash assumption is a good approximate description of the pricing-conduct of rival stores within a one-week time frame. However, conduct is much more competitive in the long run as $\omega = 0.03$, which means that every \$1.00 increase in the price index for store 1 can be expected to be matched by a \$0.03 price increase by store 2, and vice versa, in a Markov-Perfect Nash equilibrium (MPE). Prices are strategic complements in both the long- and short-runs as increases (or reductions) in the general level of prices at one store are matched with higher (or lower) prices in the rival. Pricing is highly competitive in the long-run, however, as prices are only weak complements: Price increases by one store are not nearly matched by higher prices in the other store in the long-run equilibrium.

[10] The first stage regression is also the regression summarized in Table 2.

[11] Other indices of energy, packaging, distribution, and product-content prices where not significant so were excluded from the final model. Retail wages were retained because of their *a priori* importance to retail pricing.

[table 5 in here]

Conditioning on how rivals set prices in MPE is necessary to examine the effect of switching costs on equilibrium prices. In order to calculate the effect of switching costs, and the real-option component of switching costs, on equilibrium prices, we conduct the counterfactual simulations described above. In table 6, the "With Loyalty" simulation serves as the base case, which is the set of equilibrium prices calculated under the parameter scenario shown in tables $4-6$ above. The "No Loyalty" scenario involves calculating equilibrium prices without the loyalty effect influencing each store-share, and the equilibrium prices reported in table 5. Finally, the "No Real Option" scenario involves re-simulating the set of equilibrium prices with the real option effect removed from the loyalty variable according to the marginal effects reported in table 5. According to the results reported in table 6, equilibrium prices are significantly higher without the loyalty effect included. In other words, loyalty is pro-competitive as the retailers compete more aggressively for locked-in consumers (Dube, Hitsch, and Rossi 2009), or are simply compensating them for not switching to a rival store (Arie and Grieco 2014). If the price-lowering effect of loyalty is indeed best explained by the compensation model of Arie and Grieco (2014), then some of this compensation must be for the real option embedded in observed switching costs. Comparing the implicit values of the real-option effects between stores, we find that the option value differs considerably, and accounts for a substantial proportion of the estimated switching costs. In fact, as a percentage of estimated switching costs, the real option component ranges from 57% for store 1, to 48% for consumers in store 2. Although there are many factors that enter into real option values, this difference is likely due to more volatile prices in store 2 relative to store 1 (19.3% versus 17.2%). Regardless of the relative effects, our findings show that if stores want to price in order to attract loyal consumers, they have to reduce prices to do so.

[table 6 in here]

Our findings have many aimportant implications, both for the assumed efficiency of retail pricing, and managerial conduct. First, it is generally assumed that increasing concentration among supermarket operators in the US represents the outcome of a longer-term search to soften the notoriously competitive environment among US grocery retailers (USDA 2017). As retailers become larger and develop more sophisticated loyalty programs across multiple bricks-and-mortar footprints, and across online and offline channels, our findings suggest that they will become more competitive, and retail price deflation will continue. Instead of a bulwark against competitive pricing, loyalty programs may be a facilitating tool. Second, the importance of the real-option effect in supporting customer loyalty suggests that Hi-Lo stores (stores that maintain relatively high shelf prices and compete on periodic promotions) may have an implicit advantage in building loyalty relative to Everyday-Low-Price (EDLP) stores. Because Hi-Lo customers are uncertain as to when their favorite products will be offered on sale, competitors have to offer greater discounts in order to pay for the real option built into Hi-Lo shoppers store-choice decisions. Third, retail managers are likely not directly aware of the effect loyalty programs have on retail pricing. While they may focus on intermediate goals such as shopper-retention rates and floor-traffic, only the most sophisticated retailers are fully aware of the implicit cost of achieving these goals on bottom-line profit.

5 Conclusions and Implications

In this paper, we examine the question of whether switching-costs, specifically the cost of switching loyalties between retail stores, are pro-competitive or anti-competitive. While orthodox economic theory maintains that if the cost of switching from one option to another are large, so prices will be higher as sellers take advantage of their customers' reluctance to switch to a competitor (Klemperer 1987), more recent research argues for the opposite: Namely that switching costs can be pro-competitive either due to the dominance of the investment-effect over the harvesting-effect (Dube, Hitsch and Rossi 2009; Rhodes 2014a) or due to retailers compensating switchers for their effort (Arie and Grieco 2014). We add another explanation for the possible pro-competitive effect of switching costs in that retailers must pay customers to exercise their real option to switch, and test this theory using an empirical model of store-switching behavior.

Our empirical model of store-switching behavior assumes that switching costs arise due to a "mere loyalty" effect, or the behavioral tendency to choose a familiar option in order to minimize the costs of search, and a "real option" effect. Consumers' shopping behavior gives rise to a real option effect because store-search involves fixed costs, retail prices are inherently uncertain, and consumers have a unique opportunity to switch, or to stay, at each shopping occasion. Because of this real option effect, rival stores have to "pay" to attract anothers' loyal customers through lower retail prices. The result is a pro-competitive effect due to rival retailers bidding away each others' loyal customers' real options.

We estimate a dynamic structural model of competitive retail pricing to uncover Markov Perfect pricing strategies, conditional on an intertemporal model of retail demand. Our demand model is a nested-logit specification of purchase incidence and store choice, which we estimate using data from an environment in which there are essentially only two competing stores. By allowing store choice to depend critically on the extent of loyalty exhibited toward either store, and loyalty to depend on the volatility of prices offered in each store, we are able to estimate the importance of both loyalty and the real option effect on prices in a dynamic equilibrium.

We find that loyalty is the most important variable in driving store choice among our sample of shoppers, and that price-volatility is positively related to the extent of loyalty. We interpret this as the hysteretic effect associated with real-option pricing on decisions regarding long-lived assets, such as the stream of household-shopping behavior. Using counterfactual simulations of the Markov Perfect pricing equilibrium under conditions of observed loyalty and no loyalty, we find that prices are, on average, 4.1% lower when customers exhibit loyalty to one store or another, and have to be compensated to switch. Further, we find that 52.5% of this loyalty effect is attributable to the existence of a real option in store choice, and 47.5% to "mere loyalty". Therefore, we find that loyalty is pro-competitive, in contrast to the received wisdom on this topic.

Our findings are of importance both due to the economic significance of the retail sector of the US economy, and the aggressiveness of retailers' investment in loyalty-generation programs. Retail managers, presumably, invest in loyalty programs on the assumption that they are profit-increasing, but our findings suggest that they are instead creating a more competitive environment that can be almost completely counter-productive. While the outcome of retailers' investments may be welfare-enhancing for the economy as a whole, it also implies a waste of resources, and misdirected marketing efforts. If loyalty programs are misguided, however, then our findings beg the question of what an optimal retail strategy should look like? Perhaps cherry-picking profitable shoppers with deep single-item discounts or advertisements would be the logical alternative to loyalty. Rather than encourage repeat shoppers, retailers should instead focus on the profitable "non-aligned" segment of the market that is inherently flexible, and willing to search for what looks like a deal.

6 Appendix: Real Option Values and Switching

* to here - change to be more consistent with the new empirical model. Recall, in Section 2, we defined the difference between indices of shopping-basket prices in two competing retail stores, r and s, as $g_t = p_{st} - p_{rt}$ for period t and the relationship between the two price indices as following a Geometric Brownian Motion process:

$$dg = \mu g dt + \sigma g dz, \tag{14}$$

suppressing the time subscript, where μ is the mean drift rate, σ is the standard deviation of the process, and dz defines the Wiener increment with properties: $E(dz) = 0, E(dz^2) = dt$.

As explained in Section 2, faced with a switching decision at each period, the maximum value a household can attain occurs when it purchases from the store with the lowest total cost, assuming r is the current store, or:

$$V^0(g, p_{rt}, t) = \max[V^r(g, p_{rt}), V^s(p_{st}) - \lambda c_t], \tag{15}$$

where V^r, V^s are values of a household that purchases from store r and s, respectively, λ is the marginal cost of switching, c_t, assuming the cost is only incurred when a switch takes place. The solution to this problem provides upper and lower threshold values of g_t that constitute optimal switching points between responding and not responding to a particular value of the price gap between the two stores.

The problem described in (15) above can be solved using dynamic programming techniques similar to those used to solve the learning model developed in Erdem and Keane (1996). However, Dixit (1989, 1992) develops an equivalent approach grounded in contingent claims analysis that is more consistent with the notion of a real option – that household loyalty to one store, in essence, is a claim on an asset that must be purchased by the retailer if the household is induced to switch. In this framework, the value of waiting to switch stores is a contingent security where the value of the contingency depends on the underlying price-process upon which the decision depends.

Compare the value of a household that decides to switch with one that does not using the value functions in (15). With switching costs, and uncertainty of the form shown in (1), the value of a household that does not switch consists of one capitalized stream of purchase decisions, plus the value of the option to switch to the other store, while the value of a household that does switch is the discounted value of a consumption stream that includes the most current decision. Given the process for g_t shown in (1), the equilibrium condition for a household that decides not to switch is found by first equating the instantaneous expected return on household wealth due to a change in the price gap with the required return on the household's wealth (ρV^r). In this case, the instantaneous expected return consists only of the expected rise in the value of the household (Dixit 1989) so the partial differential equation becomes:

$$(1/2)\sigma^2 g^2 V_{gg}^r + \mu g V_g^r - \rho V^r = 0, \tag{16}$$

where we note that subscripts indicate partial differentiation, we suppress the household indices for clarity, ρ is the required rate of return, and the other variables are as previously defined. The general solution to (16) is then found by substituting $g^{\theta} = V^r$ and solving the resulting quadratic equation:

$$\phi(\theta) = (1/2)\sigma^2\theta(\theta - 1) + \mu\theta - \rho = 0, \tag{17}$$

where θ is a constant to be determined. This equation has two roots: $\beta_1 < 0$, and $\beta_2 > 1$ such that:

$$\beta_1, \beta_2 = 1/2\left(\frac{1-2\mu}{\sigma^2} + / - \left(\left(1 - \frac{2\mu}{\sigma^2}\right)^2 + \frac{8\rho}{\sigma^2}\right)^{1/2}\right). \tag{18}$$

the general solution to the differential equation (16) for a household that does not switch must reflect this fact so we write:

$$V^r = A^r g^{\beta_1} + B^r g^{\beta_2}, \tag{19}$$

which must then be solved for values of A^r and B^r. The fact that this expression solves (16) can be verified by taking the first and second derivatives of V^r in (19) and substituting the result into (16). As in Dixit (1989), the two terms on the right-side of (19) represent the option to switch, while the value of not switching is normalized to zero by assumption.

We determine the value of the option by applying the value-matching and smooth pasting conditions to the expressions for the value of a non-switching household, and one that responds to the price differential by switching stores. The value of a household that decides to switch is determined solely by the flow of services provided by the basket of goods at the other store as they have already exercised the option to switch. The value of the good, in turn, is assumed to be equal to the price differential as it represents the difference between what a household is willing to pay for the basket of goods at the new store, and what the market requires them to pay. Because the price gap is assumed to drift at an average rate of μ, the present value of this stream of benefits to the household is found as:

$$V^{\alpha} = \frac{g}{\rho - \mu}, \tag{20}$$

for a constant α, so the value of the household is comprised entirely of the value of purchasing the basket of goods at the other store. Solving for the constant terms is possible if we recognize that for small values of g the value of the option to wait will be very low, so $A^r = 0$. Further, at a sufficiently high level of g, call it the upper-threshold, or g^U, a household will immediately switch so the value of a non-switching household must equal the value of a household that switches less the fixed cost of search:

$$V^r(g^U) = B^r(g^U)^{\beta_2} = \frac{g^U + p - rp}{\rho - \mu} - \lambda c = V^a(g^U), \tag{21}$$

because the value of a household that purchases is equal to the capitalized value of all desired consumption. Next, the smooth pasting condition requires that the incremental value of waiting to switch as g rises must be equal to the incremental value of a household that has already purchased, or:

$$V_g^n(g^U) = \beta_2 B^n(g^U)^{\beta_2 - 1} = 1/(\rho - \mu) = V_g^a(g^U). \tag{22}$$

The smooth pasting and value matching conditions yield two equations and two unknowns. Solving these for g^U gives a closed-form expression for the threshold price gap as a function of the parameters of the model:

$$g^U = \psi \left(\lambda c(\rho - \mu) - \rho/\mu\right), \tag{23}$$

where $\psi = \left(\frac{\beta_2}{\beta_2 - 1}\right)$ and β_2 is a function the drift rate and volatility of the price gap, and is strictly greater than one. The price gap is the difference in prices that would induce a consumer to switch from one firm to another. Because the expression on the right-side is unambiguously positive for any reasonable parameterization, this derivation shows that the presence of a real option value in store-switching behavior causes consumers to switch at larger price-differentials than would otherwise be the case. Expressed in terms of the rate of store-switching, our real option model implies a hysteretic effect, or a substantial delay, relative to the case where a real option is either not present, or is ignored.

For our econometric model of switching behavior, the critical implication that arises from this solution follows from differentiating the solution for g^U with respect to the volatility of the price gap. Because $dg^U/d\sigma > 0$, the price gap is expected to rise in the expected volatility of prices, so the hysteretic effect in store choice, manifest in the absence of store-switching (or loyalty), rises in expected price-gap volatility. We formally test this hypothesis in our econometric model. * to here - this was added 4/23/18

References

[1] Arie, G., & Grieco, P. L. (2014). Who pays for switching costs? *Quantitative Marketing and Economics*, 12(4), 379-419.

[2] Bell, D. R., and J. M. Lattin. (1998). Shopping behavior and consumer preference for store price format: Why "large basket" shoppers prefer EDLP. Marketing Science 17, 66-88.

[3] Bell, D. R., Ho, T. H., & Tang, C. S. (1998). Determining where to shop: Fixed and variable costs of shopping. *Journal of Marketing Research*, 35, 352-369.

[4] Betancourt, R., & Gautschi, D. (1990). Demand complementarities, household production, and retail assortments. *Marketing Science*, 9(2), 146-161.

[5] Briesch, R. A., Chintagunta, P. K., & Fox, E. J. (2009). How does assortment affect grocery store choice? *Journal of Marketing Research*, 46(2), 176-189.

[6] Bronnenberg, B. J., Kruger, M. W., & Mela, C. F. (2008). Database paper—The IRI marketing data set. *Marketing Science*, 27(4), 745-748.

[7] Bucklin, R. E., & Lattin, J. M. (1992). A model of product category competition among grocery retailers. *Journal of Retailing*, 68(3), 271.

[8] Caballero, R. J., & Engel, E. M. (1999). Explaining investment dynamics in US manufacturing: a generalized (S, s) approach. *Econometrica*, 67(4), 783-826.

[9] Cabral, L. (2009). Small switching costs lead to lower prices. *Journal of Marketing Research,* 46 (August), 449-51.

[10] Cabral, L. (2016). Dynamic pricing in customer markets with switching costs. *Review of Economic Dynamics*, 20, 43-62.

[11] Ching, A. T., Erdem, T., & Keane, M. P. (2014). A simple method to estimate the roles of learning, inventories and category consideration in consumer choice. *Journal of Choice Modelling*, 13, 60-72.

[12] Davidson, R., & MacKinnon, J. G. (2004). *Econometric Theory and Methods* (Vol. 5). New York: Oxford University Press.

[13] De los Santos, B., Hortaçsu, A., & Wildenbeest, M. R. (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *American Economic Review*, 102(6), 2955-2980.

[14] Dickson, P. R., & Sawyer, A. G. (1990). The price knowledge and search of supermarket shoppers. *Journal of Marketing*, 42-53.

[15] Dixit, A. (1989). Entry and exit decisions under uncertainty. *Journal of Political Economy*, 97(3), 620-638.

[16] Dixit, A. (1992). Investment and hysteresis. *Journal of Economic Perspectives*, 6(1), 107-132.

[17] Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton University Press.

[18] Doganoglu, T. (2010). Switching costs, experience goods and dynamic price competition. *Quantitative Marketing and Economics*, 8(2), 167-205.

[19] Dube, J.-P., Hitsch, G.J., Rossi, P.E. (2009). Do switching costs make markets less competitive? *Journal of Marketing Research*, 46(5), 435–45.

[20] Dube, J.-P., Hitsch, G.J., Rossi, P.E. (2010). State dependence and alternative explanations for consumer inertia. *RAND Journal of Economics*, 41(3), 417–445.

[21] Eichenbaum, M., Jaimovich, N., & Rebelo, S. (2011). Reference prices, costs, and nominal rigidities. *American Economic Review*, 101(1), 234-262.

[22] Farrell, J. and P. Klemperer (2007). Coordination and lock-in: Competition with switching costs and network effects. in *Handbook of Industrial Organization, Vol. 3*, M. Armstrong and R. Porter, eds. Amsterdam: North-Holland, 1967-2056.

[23] Fox, E. J., & Hoch, S. J. (2005). Cherry-picking. *Journal of Marketing*, 69(1), 46-62.

[24] Geweke, J., & Keane, M. (2000). Bayesian inference for dynamic discrete choice models without the need for dynamic programming. *Simulation-Based Inference in Econometrics: Methods and Applications*, eds. RS Mariano, T. Schuermann and M. Weeks, Cambridge: Cambridge University Press, 100-131.

[25] Geweke, J., & Keane, M. (2001). Computationally intensive methods for integration in econometrics. *Handbook of Econometrics* 5, 3463-3568. Elsevier.

[26] Gómez, M. I., Rao, V. R., & McLaughlin, E. W. (2007). Empirical analysis of budget and allocation of trade promotions in the US supermarket industry. *Journal of Marketing Research*, 44(3), 410-424.

[27] Gopinath, G., Gourinchas, P. O., Hsieh, C. T., & Li, N. (2011). International prices, costs, and markup differences. *American Economic Review*, 101(6), 2450-2486.

[28] Hartmann, W. R., & Viard, V. B. (2008). Do frequency reward programs create switching costs? A dynamic structural analysis of demand in a reward program. *Quantitative Marketing and Economics*, 6(2), 109-137.

[29] Hoch, S. J., Bradlow, E. T., & Wansink, B. (1999). The variety of an assortment. *Marketing Science*, 18(4), 527-546.

[30] Jones, M. A., Mothersbaugh, D. L., & Beatty, S. E. (2000). Switching barriers and repurchase intentions in services. *Journal of Retailing,* 76(2), 259-274.

[31] Keane, M. P. (1997). Modeling heterogeneity and state dependence in consumer choice behavior. *Journal of Business & Economic Statistics*, 15(3), 310–327.

[32] Klemperer, P. (1987). The competitiveness of markets with switching costs. *RAND Journal of Economics*, 18(1), 138-150.

[33] Leszczyc, P. T. P., Sinha, A., & Timmermans, H. J. (2000). Consumer store choice dynamics: an analysis of the competitive market structure for grocery stores. *Journal of Retailing*, 76(3), 323-345.

[34] Liu, X., Derdenger, T., & Sun, B. (2017), An Empirical Analysis of Consumer Purchase Behavior of Base Products and Add-ons. *Forthcoming, Marketing Science.*

[35] Mehta, N., Rajiv, S., & Srinivasan, K. (2003). Price uncertainty and consumer search: A structural model of consideration set formation. *Marketing Science*, 22(1), 58-84.

[36] Murray, K. B., & Häubl, G. (2007). Explaining cognitive lock-in: The role of skill-based habits of use in consumer choice. *Journal of Consumer Research*, 34(1), 77-88.

[37] Osborne, M. (2011). Consumer learning, switching costs, and heterogeneity: A structural examination. *Quantitative Marketing and Economics*, 9(1), 25-70.

[38] Pavlidis, P., & Ellickson, P. B. (2015). Switching costs and market power under umbrella branding. Working paper, Simon School of Business, University of Rochester, Rochester, NY.

[39] Pearcy, J. (2015). Bargains followed by bargains: when switching costs make markets more competitive. *Journal of Economics & Management Strategy,* forthcoming.

[40] Perloff, J. M., Karp, L. S., & Golan, A. (2007). *Estimating market power and strategies.* Cambridge University Press.

[41] Petrin, A., & Train, K. (2010). A control function approach to endogeneity in consumer choice models. *Journal of Marketing Research*, 47(1), 3-13.

[42] Reichheld, F. F. (2003). The one number you need to grow. *Harvard Business Review* (https://hbr.org/2003/12/the-one-number-you-need-to-grow). Accessed on April 23, 2017.

[43] Rhee, H., & Bell D. R. (2002). The inter-store mobility of supermarket shoppers. *Journal of Retailing*, 78(4), 225-237.

[44] Rhodes, A. (2014a). Re-examining the effects of switching costs. *Economic Theory*, 57(1), 161-194.

[45] Rhodes, A. (2014b). Multiproduct retailing. *Review of Economic Studies*, 82(1), 360-390.

[46] Richards, T. J., Gómez, M. I., & Printezis, I. (2015). Hysteresis, price acceptance, and reference prices. *American Journal of Agricultural Economics*, 98(3), 679-706.

[47] Richards, T. J. and S. F. Hamilton. (2016). Retail market power in a shopping-basket model of retail competition. Working paper, W. P. Carey School of Business, Arizona State University, Tempe, AZ.

[48] Roberts, M. J., & Samuelson, L. (1988). An empirical analysis of dynamic, nonprice competition in an oligopolistic industry. *RAND Journal of Economics*, 200-220.

[49] Smith, H. (2004). Supermarket choice and supermarket competition in market equilibrium. *Review of Economic Studies*, 71(1), 235-263.

[50] Shin, J., Sudhir, K., Cabral, L.M.B., Dube, J.-P., Hitsch, G.J., Rossi, P.E. (2009). Commentaries and rejoinder to do switching costs make markets less competitive. *Journal of Marketing Research*, 46(4), 446–552.

[51] Shy, O. (2002). A quick-and-easy method for estimating switching costs. *International Journal of Industrial Organization*, 20(1), 71–87.

[52] Slade, M. E. (1999). Sticky prices in a dynamic oligopoly: An investigation of (s, S) thresholds. *International Journal of Industrial Organization*, 17(4), 477-511.

[53] Staiger, D., & Stock, J. H. (1997). Instrumental variables regression with weak instruments. *Econometrica*, 65(3), 557-586.

[54] Train, K. (2003). *Discrete choice methods with simulation*. Cambridge: Cambridge University Press.

[55] United States Department of Agriculture. Economic Research Service. 2017. Retail Trends. (https://www.ers.usda.gov/topics/food-markets-prices/retailing-wholesaling/retail-trends/). Accessed on April 25, 2017.

[56] Viard, B. (2007). Do switching costs make markets more or less competitive? The case of 800-number portability. *RAND Journal of Economics*, 38 (1), 146-63.

[57] Villas-Boas, M. (2015). A short survey on switching costs and dynamic competition. Working Paper, Hass School of Business, University of California, Berkeley, Berkeley, CA.

[58] Wildenbeest, M. R. (2011). An empirical model of search with vertically differentiated products. *RAND Journal of Economics*, 42(4), 729-757.

[59] Zauberman, G. (2003). The intertemporal dynamics of consumer lock-in. *Journal of Consumer Research*, 30(3), 405-419.

Table 1. Store-Level Data Summary

		Store 1				Store 2			
Variable	Measure	Mean	Std. Dev.	Min.	Max.	Mean	Std. Dev.	Min.	Max.
Store Share	Share	0.2639	0.2585	0.0000	0.9831	0.2106	0.2785	0.0000	1.0000
Price Index	$ / oz	4.1741	0.2926	3.5417	5.4279	3.9130	0.3113	3.1215	5.4048
Variety Index	Ave. #	397.6413	45.0237	293.2440	587.0000	306.4663	22.9816	230.5000	379.4167
Distance	Miles	1.4399	0.2957	0.8447	1.7123	1.3990	1.8888	14.9813	0.7556
Milk Volume	Pints	2.9219	6.3504	0.0000	80.0000	1.4062	3.4613	0.0000	48.0000
Cereal Volume	Pounds	0.1626	0.6590	0.0000	9.8501	0.1421	0.6237	0.0000	15.7500
Soda Volume	Gallons	0.1830	0.5700	0.0000	8.5417	0.1775	0.6643	0.0000	12.0000
Snacks Volume	Pounds	0.3081	0.9046	0.0000	13.3752	0.1584	0.4824	0.0000	8.2500
Frozen Dinners Volume	Pounds	0.0868	0.5440	0.0000	12.0000	0.1246	0.5963	0.0000	12.0000
Soup Volume	Pounds	0.1862	0.9498	0.0000	27.5757	0.1607	0.8537	0.0000	28.3067
Toilet Paper Volume	Rolls	0.7889	4.1921	0.0000	96.0000	0.6342	2.8293	0.0000	60.0000
Yogurt Volume	Pints	0.1276	0.5914	0.0000	11.6250	0.1312	0.6071	0.0000	18.7500
Beer Volume	Cases (24)	0.0533	0.2189	0.0000	4.5000	0.0710	0.2810	0.0000	5.0000
Coffee Volume	Pounds	0.2292	1.2938	0.0000	24.9500	0.1435	0.8082	0.0000	27.8741
Milk Price	$ / pint	0.3856	0.1105	0.1100	2.2950	0.4586	0.0831	0.2000	1.9867
Cereal Price	$ / pound	3.0627	0.4842	0.7920	9.4883	3.0895	0.4845	0.6400	11.0570
Soda Price	$ / gallon	3.2412	0.9899	0.8331	14.2263	3.4305	0.7206	1.0982	18.0822
Snacks Price	$ / pound	3.7785	0.7412	1.2307	15.9467	3.7624	0.6851	0.8000	11.4260
Frozen Dinner Price	$ / pound	3.4466	0.4786	0.9231	9.0434	3.2846	0.6034	1.0817	10.6758
Soup Price	$ / pound	2.3880	0.8374	0.4816	28.7900	2.2618	0.7816	0.5517	26.0174
Toilet Paper Price	$ / roll	0.6006	0.1163	0.1683	1.4817	0.5387	0.0922	0.1875	1.3817
Yogurt Price	$ / pint	1.8953	0.2766	0.5625	4.7733	1.7984	0.2625	0.5175	3.8938
Beer Price	$ / Case (24)	17.7812	1.9476	9.0320	56.8567	16.0066	2.2474	5.9900	90.5896
Coffee Price	$ / Pound	2.8507	0.5083	1.2530	9.7481	3.0857	0.5630	1.1300	9.3067

Note: Beer cases are 24-12oz bottles or cans.

Table 2. Reduced-form Loyalty Probit Model

	Store 1			Store 2		
Variable	Estimate	SE	Marg.	Estimate	SE	Marg.
Constant	-3.3103*	0.0549		-0.8006*	0.0799	
Price Ratio	0.1577*	0.0055	0.0328	0.2839*	0.0564	0.0432
Price Volatility	0.6038*	0.1842	0.1257	0.3591	0.2273	0.0546
Promotion	0.3268*	0.0460	0.0680	0.0894	0.0640	0.0136
Feature	0.0120	0.0482	0.0025	-0.2715*	0.0454	-0.0413
Income	4.6507*	0.1384	0.9683	-2.6599*	0.1788	-0.4047
Family Size	6.6354*	0.4382	1.3815	-12.2977*	0.5305	-1.8709
Age	17.5200*	0.5576	3.6476	-4.4400*	0.5925	-0.6755
Education	35.4466*	2.1062	7.3798	-16.8763*	1.9980	-2.5674
LLF	-50,570.03			-36,057.09		
χ^2	3,695.66			1,361.05		

Note: A single asterisk indicates significance at a 5% level. Marginal effects calculated at sample means. Price volatility is 4-period moving standard deviation.

Table 3. Nested Logit Purchase-Incidence Model

	Store 1		Store 2		Store 1		Store 2	
	Model 1		Model 1		Model 2		Model 2	
	Fixed Coefficients		Fixed Coefficients		Random Coefficients		Random Coefficients	
	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
Milk	-0.8521*	0.2380	0.0173	0.2157	-0.8499*	0.2381	0.0170	0.2157
Cereal	-0.7753*	0.2375	0.4241*	0.2154	-0.7732*	0.2375	0.4238*	0.2155
Soda	-0.1386	0.2369	0.9944*	0.2151	-0.1364	0.2370	0.9942*	0.2151
Snacks	-0.1471	0.2372	1.0547*	0.2152	-0.1449	0.2373	1.0544*	0.2152
Frozen Dinners	-1.6795*	0.2374	0.0518	0.2152	-1.6776*	0.2374	0.0515	0.2152
Soup	-0.9570*	0.2374	0.1983	0.2157	-0.9547*	0.2374	0.1980	0.2158
Toilet Paper	-2.3830*	0.2382	-1.1279*	0.2164	-2.3809*	0.2383	-1.1280*	0.2164
Yogurt	-1.0848*	0.2377	0.0435	0.2157	-1.0825*	0.2378	0.0431	0.2157
Beer	-1.0365*	0.2373	0.4095	0.2154	-1.0341*	0.2373	0.4091	0.2154
Lagged Volume	0.1243*	0.0007	0.1707*	0.0010	0.1243*	0.0007	0.1708*	0.0010
					0.0016*	0.0005	-0.0019*	0.0007
Consumption Rate	0.0299*	0.0016	0.0259*	0.0016	0.0299*	0.0016	0.0258*	0.0016
					-0.0009	0.0008	0.0009	0.0010
Interpurchase Time	0.2043*	0.0006	0.2275*	0.0009	0.2044*	0.0006	0.2274*	0.0009
					-0.0051*	0.0005	-0.0021*	0.0008
Income	0.5226	0.6042	-3.1295*	0.8364	0.5240	0.6044	-3.1307*	0.8366
Family Size	0.1507	1.9528	-5.0963*	2.2099	0.1361	1.9529	-5.1008*	2.2106
Age	18.6636*	2.4325	5.0290	2.6828	18.6442*	2.4329	5.0320	2.6831
Education	-8.5033	10.3260	-32.3723*	8.6441	-8.5838	10.3283	-32.3444*	8.6454
LLF	-1825088.2				-1825036.5			
AIC/N	28.5623				28.5614			

Note: A single asterisk indicates significance at a 5% level.

Table 4. Nested Logit Store-Choice Model

	Store Choice			
	Model 1		Model 2	
	Fixed Coeffs.		Random Coeffs.	
	Estimate	SE	Estimate	SE
Store 1 Constant	-2.5531*	0.0797	-2.5526*	0.0797
			-0.0132*	0.0082
Store 2 Constant	-2.2675*	0.0945	-2.2675*	0.0945
			-0.0148*	0.0097
Price	-0.1644*	0.0208	-0.1642*	0.0208
			-0.0005	0.0026
Variety	0.0053*	0.0002	0.0053*	0.0002
			-0.0002	0.0002
Loyalty	3.5086*	0.0372	3.5092*	0.0372
			-0.0233	0.0251
Distance	-0.0164*	0.0040	-0.0164*	0.0040
			0.0009	0.0031
Promotion	-0.2116*	0.0768	-0.2118*	0.0768
Feature	0.2605*	0.0399	0.2608*	0.0399
Control Function	0.5226*	0.0462	0.5228*	0.0462
λ	0.0351*	0.0074	0.0349*	0.0074

Note: Goodness of fit metrics in table 3. A single asterisk indicates ignificance at a 5% level. SE = Standard error.

Table 5. Optimal Intertemporal Pricing Model

	Model 1		Model 2	
	Fixed Coeffs.		Random Coeffs.	
	Estimate	SE	Estimate	SE
Store 1	-5.6489*	0.5372	-5.6172*	0.5525
Store 2	-5.9607*	0.5371	-5.9180*	0.5523
Labor Cost	-0.0289*	0.0110	-0.0247*	0.0114
Business Services	0.5242*	0.0242	0.5162*	0.0249
ϕ	0.9485*	0.0009	0.6389*	0.0050
$\phi(s)$			0.0826*	0.0031
ω	0.0323*	0.0070	0.1207*	0.0557
$\omega(s)$			-0.0156*	0.0541
LLF	-175.1117		-195.6481	
AIC	0.5919		0.6590	

Note: A single asterisk indicates significance at a 5% level. NL IV estimation method used.

Table 6. Pricing Simulations

		Price	SD	t-ratio
Store 1	With Loyalty	4.1858	0.2423	
	No Loyalty	4.3302*	0.1784	11.9030
	No Real Option	4.2480*	0.1698	5.2100
Store 2	With Loyalty	3.9149	0.3604	
	No Loyalty	4.0936*	0.1916	10.8556
	No Real Option	4.0074*	0.1999	5.5644

Note: A single asterisk indicates significance at a 5% level.